




# LANPLEX® 2500 EXTENDED SWITCHING USER GUIDE

Part No. 801-00343-000  
Published November 1996  
Revision 02



**3Com Corporation ■ 5400 Bayfront Plaza ■ Santa Clara, California ■ 95052-8145**

© 3Com Corporation, 1996. All rights reserved. No part of this documentation may be reproduced in any form or by any means or used to make any derivative work (such as translation, transformation, or adaptation) without permission from 3Com Corporation.

3Com Corporation reserves the right to revise this documentation and to make changes in content from time to time without obligation on the part of 3Com Corporation to provide notification of such revision or change.

3Com Corporation provides this documentation without warranty of any kind, either implied or expressed, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. 3Com may make improvements or changes in the product(s) and/or the program(s) described in this documentation at any time.

**UNITED STATES GOVERNMENT LEGENDS:**

If you are a United States government agency, then this documentation and the software described herein are provided to you subject to the following restricted rights:

**For units of the Department of Defense:**

*Restricted Rights Legend:* Use, duplication or disclosure by the Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) for restricted Rights in Technical Data and Computer Software clause at 48 C.F.R. 52.227-7013. 3Com Corporation, 5400 Bayfront Plaza, Santa Clara, California 95052-8145.

**For civilian agencies:**

*Restricted Rights Legend:* Use, reproduction or disclosure is subject to restrictions set forth in subparagraph (a) through (d) of the Commercial Computer Software - Restricted Rights Clause at 48 C.F.R. 52.227-19 and the limitations set forth in 3Com's standard commercial agreement for the software. Unpublished rights reserved under the copyright laws of the United States.

3ComFacts, Ask3Com, CardFacts, NetFacts, and CardBoard are service marks of 3Com Corporation.

3Com, LANplex, Transcend, and NETBuilder II are registered trademarks of 3Com Corporation.

CompuServe is a registered trademark of CompuServe, Inc.

3Com registered trademarks are registered in the United States, and may or may not be registered in other countries.

Other brand and product names may be registered trademarks or trademarks of their respective holders.

Guide written, edited, and illustrated by Trish Crawford, Lynne Gelfand, Michael Jenness, Dave Sullivan, Patricia Johnson, Michael Taillon, Iain Young, and Bonnie Jo Collins.

# CONTENTS

---

## ABOUT THIS GUIDE

- Introduction 1
- How to Use This Guide 1
- Conventions 2
- LANplex 2500 Documentation 3
- Documentation Comments 5

## PART I GETTING STARTED

---

### 1 LANPLEX® EXTENDED SWITCHING FEATURES

- About LANplex Extended Switching 1-1
- Using Menus 1-2
  - Bridge Menu 1-3
  - IP Menu 1-4
  - IPX Menu 1-5
  - Appletalk Menu 1-6

## PART II VIRTUAL LAN TECHNOLOGY

---

### 2 VLANS ON THE LANPLEX® SYSTEM

- About VLANs 2-1
  - Types of VLANs 2-1
    - Port Group VLANs 2-1
    - MAC Address Group VLANs 2-2
    - Application-Oriented VLANs 2-2
    - Protocol-Sensitive VLANs 2-2
  - LANplex Protocol-Sensitive VLAN Configuration 2-3
    - Protocol Suite 2-3
    - Switch Ports 2-4
    - Layer 3 Addressing Information 2-4
  - Default VLAN 2-5

|  |      |
|--|------|
| Modifying the Default VLAN                       | 2-5  |
| How the LANplex® System Makes Flooding Decisions | 2-5  |
| VLAN Exception Flooding                          | 2-6  |
| Overlapped IP VLANs                              | 2-7  |
| Routing Between VLANs                            | 2-8  |
| VLAN Examples                                    | 2-10 |
| Example 1  | 2-10 |
| Example 2  | 2-11 |

## **PART III    ABOUT ROUTING PROTOCOLS**

---

### **3    BRIDGING AND ROUTING IN THE LANPLEX® SYSTEM**

|  |     |
|--|-----|
| What Is Routing?                       | 3-1 |
| LANplex in a Subnetworked Environment  | 3-2 |
| Integrating Bridging and Routing       | 3-3 |
| Bridging and Routing Models            | 3-4 |
| Traditional Bridging and Routing Model | 3-4 |
| LANplex Bridging and Routing Model     | 3-6 |

---

### **4    ROUTING WITH IP TECHNOLOGY**

|                                    |      |
|------------------------------------|------|
| IP Routing and the OSI Model       | 4-1  |
| Elements of IP Routing             | 4-2  |
| IP Addresses                       | 4-2  |
| Address Classes                    | 4-3  |
| Subnet Part of an IP Address       | 4-3  |
| Router Interfaces                  | 4-4  |
| Routing Table                      | 4-5  |
| Static Routes                      | 4-6  |
| Dynamic Routes Using RIP           | 4-6  |
| Default Route                      | 4-7  |
| Address Resolution Protocol (ARP)  | 4-7  |
| IP Routing Transmission Errors     | 4-9  |
| Routing with Classical IP over ATM | 4-10 |
| About Logical IP Subnets (LISs)    | 4-10 |
| ATM ARP Servers                    | 4-10 |
| Forwarding to Nodes within an LIS  | 4-11 |
| IP Routing References              | 4-11 |

---

## 5 ROUTING WITH IP MULTICAST

- About IP Multicast Routing 5-1
- IGMP 5-1
- DVMRP 5-2
  - The MBONE 5-2
- Multicast Routing
  - Algorithms 5-3
  - Flooding 5-3
  - Spanning Trees 5-3
  - Reverse Path Forwarding 5-4
  - Pruning 5-5
- Multicast Interfaces 5-5
  - DVMRP Metric Value 5-5
  - Time-To-Live (TTL) Threshold 5-5
  - Rate Limit 5-6
- Multicast Tunnels 5-6

---

## 6 ROUTING WITH IPX

- IPX Routing in the NetWare® Environment 6-1
  - Internet Packet Exchange (IPX) 6-2
  - Routing Information Protocol (RIP) 6-3
  - Service Advertising Protocol (SAP) 6-3
- How IPX Routing Works 6-4
  - IPX Packet Format 6-4
  - IPX Packet Delivery 6-6
    - Sending Node's Responsibility 6-6
    - Router's Responsibility 6-7
- The Elements of
  - IPX Routing 6-8
  - Router Interfaces 6-8
  - Routing Tables 6-8
    - Generating Routing Table Information 6-9
    - Selecting the Best Route 6-10
  - Service Advertising Protocol 6-10
    - Internetwork Service Information 6-10
    - SAP Packet Structure 6-11
    - Server Information Table 6-13
    - Server Information Maintenance 6-14

---

## **7 ROUTING IN AN APPLE TALK® ENVIRONMENT**

- About AppleTalk® 7-1
- AppleTalk® Network Elements 7-1
  - AppleTalk® Networks 7-2
  - AppleTalk® Nodes 7-2
    - Named Entities 7-2
  - AppleTalk® Zones 7-3
  - Seed Routers 7-4
- AppleTalk Protocols 7-4
  - Physical Connectivity 7-5
  - The Datagram Delivery Protocol (DDP) 7-6
  - End-to-End Services 7-6
    - Transport Layer Protocols 7-6
    - The Session Layer Protocols 7-9
  - Presentation Layer 7-10
- About AARP 7-10

## **PART IV ADMINISTERING EXTENDED SWITCHING FEATURES**

---

### **8 ADMINISTERING VLANs**

- Displaying VLAN Information 8-1
- Defining VLAN Information 8-3
- Modifying VLAN Information 8-4
- Removing VLAN Information 8-5

---

### **9 ADMINISTERING IP ROUTING**

- Administering interfaces 9-1
  - LIS Interfaces 9-2
  - Interface Characteristics 9-2
  - Displaying Interfaces 9-3
  - Defining an IP LIS Interface 9-4
  - Defining an IP VLAN Interface 9-6
  - Modifying an Interface 9-7
  - Removing an Interface 9-7
  - Adding an Advertisement Address 9-8
  - Removing an Advertisement Address 9-8
  - Adding a Permanent Virtual Circuit (PVC) 9-9
  - Removing a Permanent Virtual Circuit (PVC) 9-9
- Administering Routes 9-9
  - Displaying the Routing Table 9-11

|  |      |
|--|------|
| Defining a Static Route                      | 9-11 |
| Removing a Route                             | 9-12 |
| Flushing a Route                             | 9-12 |
| Setting the Default Route                    | 9-12 |
| Removing the Default Route                   | 9-13 |
| Administering the ARP Cache                  | 9-13 |
| Displaying the ARP Cache                     | 9-14 |
| Removing an ARP Cache Entry                  | 9-14 |
| Flushing the ARP Cache                       | 9-15 |
| Administering ATM ARP Servers                | 9-15 |
| Displaying ATM ARP Servers                   | 9-15 |
| Defining an ATM ARP Server                   | 9-16 |
| Removing an ATM ARP Server                   | 9-16 |
| Displaying the ATM ARP Cache                 | 9-17 |
| Removing an ATM ARP Cache Entry              | 9-17 |
| Flushing the ATM ARP Cache                   | 9-18 |
| Administering UDP Helper                     | 9-18 |
| Displaying UDP Helper Information            | 9-19 |
| Defining a Port and an IP Forwarding Address | 9-19 |
| Removing a Port or an IP Forwarding Address  | 9-19 |
| Setting the BOOTP Hop Count Limit            | 9-20 |
| Setting the BOOTP Relay Threshold            | 9-20 |
| Enabling and Disabling IP Routing            | 9-20 |
| Enabling and Disabling ICMP Router Discovery | 9-21 |
| Setting the RIP Mode                         | 9-21 |
| Pinging an IP Station                        | 9-22 |
| Displaying IP Statistics                     | 9-23 |

---

## 10 ADMINISTERING IP MULTICAST ROUTING

|                                       |      |
|---------------------------------------|------|
| Enabling and Disabling DVMRP          | 10-2 |
| Enabling and Disabling IGMP           | 10-2 |
| Administering IP Multicast Interfaces | 10-3 |
| DVMRP Metric Value                    | 10-3 |
| Time To Live (TTL) Threshold          | 10-3 |
| Rate Limit                            | 10-4 |
| Displaying Multicast Interfaces       | 10-4 |
| Disabling Multicast Interfaces        | 10-5 |
| Enabling Multicast Interfaces         | 10-5 |
| Administering Multicast Tunnels       | 10-6 |
| Displaying Multicast Tunnels          | 10-6 |
| Defining a Multicast Tunnel           | 10-7 |
| Removing a Multicast Tunnel           | 10-7 |

Displaying Routes 10-8  
Displaying the Multicast Cache 10-9

---

## **11 ADMINISTERING IPX ROUTING**

Administering Interfaces 11-2  
    Displaying IPX Interfaces 11-3  
    Defining an IPX Interface 11-3  
    Modifying an Interface 11-4  
    Removing an Interface 11-4  
Administering Routes 11-5  
    Displaying the Routing Table 11-6  
    Defining a Static Route 11-6  
    Removing a Route 11-7  
    Flushing Routes 11-7  
Administering Servers 11-8  
    Displaying the Server Table 11-9  
    Defining a Static Server 11-9  
    Removing a Server 11-10  
    Flushing Servers 11-10  
Setting IPX Forwarding 11-11  
Setting the RIP Mode 11-11  
Setting the Enhanced RIP Mode 11-12  
Setting the SAP Mode 11-13  
Displaying Statistics 11-14  
    Displaying IPX Summary Statistics 11-14  
    Displaying IPX RIP Statistics 11-15  
    Displaying IPX SAP Statistics 11-16  
    Displaying IPX Forwarding Statistics 11-17

---

## **12 ADMINISTERING APPLE TALK® ROUTING**

Administering Interfaces 12-2  
    Displaying AppleTalk Interfaces 12-3  
    Defining an Interface 12-3  
    Removing an Interface 12-4  
Administering Routes 12-5  
    Displaying the Routing Table 12-5  
    Flushing all Routes 12-6  
Administering the AARP Cache 12-7  
    Displaying the AARP Cache 12-8  
    Removing an Entry in the Cache 12-9  
    Flushing All Cache Entries 12-9  
Displaying the Zone Table 12-10



- Configuring Forwarding 12-11
- Configuring Checksum 12-12
- Pinging an AppleTalk Node 12-12
- Viewing Appletalk Statistics 12-13
  - Displaying DDP Statistics 12-13
  - Displaying RTMP Information 12-14
  - Displaying ZIP Information 12-15
  - Displaying NBP Information 12-17

## **PART V    REMOTE MONITORING (RMON) AND THE LANPLEX® SYSTEM**

---

### **13    REMOTE MONITORING (RMON) TECHNOLOGY**

- What Is RMON? 13-1
- Benefits of RMON 13-2
- LANplex RMON Implementation 13-2
  - 3Com Transcend RMON Agents 13-3
- Management Information Base (MIB) 13-4
  - MIB Objects 13-4
- Alarms 13-6
  - Setting Alarm Thresholds 13-7
  - Example of an Alarm Threshold 13-7
  - RMON Hysteresis Mechanism 13-8

## **PART VI    APPENDIX**

---

### **A    TECHNICAL SUPPORT**

- On-line Technical Services A-1
  - 3Com Bulletin Board Service A-1
    - Access by Analog Modem A-1
    - Access by Digital Modem A-2
  - World Wide Web Site A-2
  - 3ComForum on CompuServe® A-2
  - 3ComFacts™ Automated Fax Service A-3
- Support from Your Network Supplier A-3

Support from 3Com A-4  
Returning Products for Repair A-4

---

## INDEX

# ABOUT THIS GUIDE

---

## Introduction

The *LANplex® 2500 Extended Switching User Guide* provides information about the features included with the LANplex Extended Switching software. These features include IP, IP Multicast, classical IP over ATM, IPX, and AppleTalk routing, virtual LAN (VLAN) configuration, and remote monitoring (RMON).

Use this guide with the *LANplex® 2500 Administration Console User Guide* when you configure your LANplex 2500 system.



*See the LANplex® 2500 Software Installation and Release Notes for information about how to install Extended Switching software on your LANplex system.*

### *Audience description*

This guide is intended for the system or network administrator who is responsible for configuring, using, and managing the LANplex 2500 system. It assumes a working knowledge of local area network (LAN) operations and a familiarity with communications protocols used on interconnected LANs.



*If the information in the release notes shipped with your product differs from the information in this guide, follow the release notes.*

---

## How to Use This Guide

The following table shows where to find specific information.




| If you are looking for...                                | Turn to... |
|--|------------|
| An overview of Extended Switching features               | Chapter 1  |
| Virtual LANs (VLANs) on the LANplex System               | Chapter 2  |
| General routing and routing models in the LANplex system | Chapter 3  |
| IP routing strategies                                    | Chapter 4  |
| IP multicast routing and its protocols                   | Chapter 5  |
| continued  |            |

| If you are looking for...                          | Turn to... |
|--|------------|
| IPX routing and its protocols                      | Chapter 6  |
| AppleTalk routing, network elements, and protocols | Chapter 7  |
| How to administer VLANs                            | Chapter 8  |
| How to administer IP routing                       | Chapter 9  |
| How to administer IP mulitcast routing             | Chapter 10 |
| How to administer IPX routing                      | Chapter 11 |
| How to administer AppleTalk routing                | Chapter 12 |
| Remote Monitoring (RMON)                           | Chapter 13 |
| 3Com Technical Support                             | Appendix A |

## Conventions

Table 1 and Table 2 list conventions that are used throughout this guide.

**Table 1** Notice Icons

| Icon  | Type             | Description   |
|---|------------------|---|
|  | Information Note | Information notes call attention to important features or instructions.     |
|  | Caution          | Cautions alert you to personal safety risk, system damage, or loss of data. |
|  | Warning          | Warnings alert you to the risk of severe personal injury.                   |

**Table 2** Text Conventions

| Convention             | Description   |
|------------------------|---|
| "Enter"                | "Enter" means type something, then press the [Return] or [Enter] key.   |
| "Syntax" vs. "Command" | <p>"Syntax" indicates that the general command syntax form is provided. You must evaluate the syntax and supply the appropriate value; for example:</p> <p>Set the date by using the following syntax:</p> <pre>mm/DD/yy hh:mm:ss xm</pre> <p>"Command" indicates that all variables in the command syntax form have been supplied and you can enter the command as shown in text; for example:</p> <p>To update the system software, enter the following command:</p> <pre><b>system software Update</b></pre> |
| screen display         | <p>This <i>typeface</i> indicates text that appears on your terminal screen; for example:</p> <pre>NetLogin:</pre>  |
| <b>commands</b>        | <p><b>This typeface</b> indicates commands that you enter; for example:</p> <pre><b>bridge port stpState</b></pre>  |
| <i>Italic</i>          | <i>Italic</i> is used to denote emphasis and buttons.   |
| Keys                   | <p>When specific keys are referred to in the text, they are called out by their labels, such as "the Return key" or "the Escape key," or they may be shown as [Return] or [Esc].</p> <p>If two or more keys are to be pressed simultaneously, the keys are linked with a plus sign (+), for example:</p> <p>Press [Ctrl]+[Alt]+[Del].</p>   |

## LANplex 2500 Documentation

The following documents comprise the LANplex 2500 documentation set. If you want to order a document that you do not have or order additional documents, contact your sales representative for assistance.

- *LANplex® 2500 Unpacking Instructions*

Describe how to unpack your LANplex system. It also provides you with an inventory list of all the items shipped with your system. (Shipped with system/Part No. 801-00353-00)

- *LANplex® 2500 Software Release Notes*  
Provide information about the software release, including new features and bug fixes. It also provides information about any changes to the LANplex system's documentation. (Shipped with system)
- *LANplex® 2500 Getting Started*  
Describes all the procedures necessary for installing, cabling, powering up, configuring management access to, and troubleshooting your LANplex system. (Shipped with system/Part No. 801-00355-000)
- *LANplex® 2500 Operation Guide*  
Provides information to help you understand system management and administration, bridging, Fast Ethernet, ATM, and FDDI technology. It also describes how these concepts are implemented in the LANplex system. (Shipped with system/Part No. 801-00344-000)
- *LANplex® 2500 Administration Console User Guide*  
Provides information about using the Administration Console to configure and manage your LANplex system. (Shipped with system/Part No. 801-00322-000)
- *LANplex® 2500 Extended Switching User Guide (This book)*  
Describes® how the routing protocols, VLAN, and RMON are implemented in the LANplex system and provides information about using the Administration Console to configure and manage these features. (shipped with the option package/Part No. 801-00343-000)
- *LANplex® 2500 Intelligent Switching Administration Console Command Quick Reference card*  
Contains the Administration Console Intelligent Switching commands for the LANplex system. (Shipped with the system/Part No. 801-000318-000)
- *LANplex® 2500 Extended Switching ADMINISTRATION CONSOLE Command Quick Reference card*  
Contains the Administration Console Extended Switching commands for the LANplex system. (Shipped with the option package/Part No. 801-00319-000)

- *Module Installation Guides*

Provide an overview, installation instructions, LED status information, and pin-out information for the particular option module. (Shipped with individual modules)

---

## Documentation Comments

Your suggestions are very important to us and will help make our documentation more useful to you. Please email comments about this document to 3Com at: **sdtechpubs\_comments@3Mail.3Com.com**

Please include the following information when commenting:

- Document title
- Document part number (listed on back cover of document)
- Page number (if appropriate)

*Example:*    *LANplex® 2500 Operation Guide*  
                  Part No. 801-00344-000  
                  Page 2-5 (chapter 2, page 5)



6

ABOUT THIS GUIDE



# 1

## LANPLEX® EXTENDED SWITCHING FEATURES

This chapter provides an overview of the Extended Switching software, and describes the enhanced Administration Console menus.

---

### About LANplex Extended Switching

The LANplex Extended Switching software replaces your existing LANplex software and adds new functionality to your system. Extended Switching software contains all the features of LANplex Intelligent Switching software, in addition to:

- Virtual LANs (VLANs)
- Internet Protocol (IP) Routing (an enhanced version of IP from the standard system software)
- IP multicast routing
- Classical IP routing over Asynchronous Transfer Mode (ATM)
- Internet Packet Exchange (IPX) routing
- AppleTalk® routing
- Remote Monitoring (RMON)

For information on how to gain access to online help, to use scripts, and to exit from the Administration Console, see the *LANplex® 2500 Administration Console User Guide*.



*See the LANplex® 2500 Software Installation and Release Notes for information about how to install Extended Switching software on your LANplex system.*

## Using Menus

When you gain access to the Administration Console, the top-level menu appears. The Extended Switching software contains top-level menus and additions to the Bridge and IP menu options not available with Intelligent Switching software:

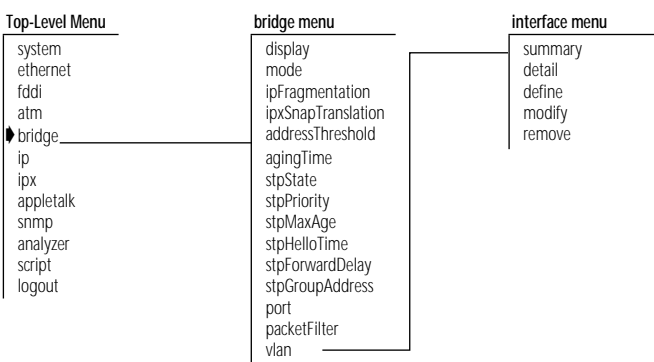
|                                      | Menu options          | Option Descriptions                    |
|--------------------------------------|-----------------------|--|
| Menu options vary by level of access | system                | - Administer system-level functions    |
|                                      | ethernet              | - Administer Ethernet ports            |
|                                      | fddi                  | - Administer FDDI resources            |
|                                      | ATM                   | - Administer ATM resources             |
|                                      | bridge                | - Administer bridging/VLANs            |
|                                      | ip                    | - Administer IP                        |
|                                      | ipx                   | - Administer IPX                       |
|                                      | appletalk             | - Administer Appletalk                 |
|                                      | snmp                  | - Administer SNMP                      |
|                                      | analyzer              | - Administer Roving Analysis           |
|                                      | script                | - Run a script of console commands     |
|                                      | logout                | - Logout of the Administration Console |
|                                      | Type ? for help.      |  |
|                                      | Select a menu option: |  |

The following sections show the enhanced menus provided with Extended Switching software. All other menu items appear in the *LANplex® 2500 Administration Console User Guide*.



*The RMON feature is available through SNMP only. This feature is not available through the Administration Console. See Chapter 13, Remote Monitoring (RMON) Technology, for more information about this feature.*

**Bridge Menu** From the **bridge** menu, you can view information about and configure Ethernet LANs, including VLANs. Figure 1-1 shows the **bridge** menu.



**Figure 1-1** Bridge Menu Hierarchy

**IP Menu** From the **ip** menu, you can view information about and configure Internet Protocol (IP) interfaces and routes as well as IP Multicast routing. You can administer the Address Resolution Protocol (ARP), the Routing Information Protocol (RIP), UDP Helper, IP Forwarding, and ping IP stations. You can also define ATM ARP servers from the **ip** menu if you are running classical IP over ATM. Figure 1-2 shows the **ip** menu. To define a new IP interface, for example, enter **ip** at the top-level menu, **interface** at the ip menu, and then **define** at the interface menu.

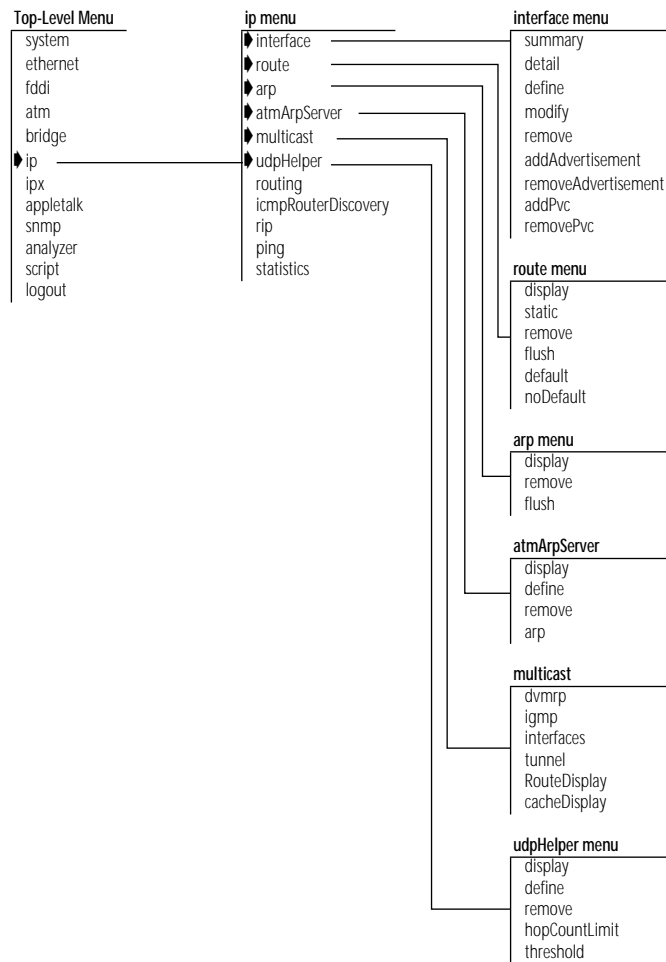


Figure 1-2 IP Menu Hierarchy

**IPX Menu** From the **ipx** menu, you can view information about and configure Internet Packet Exchange (IPX) interfaces, routes, and servers. You can also administer the Routing Information Protocol (RIP), Enhanced RIP mode, Service Advertising Protocol (SAP), and statistics. Figure 1-3 shows the IPX menu. For example, to define a new IPX interface, enter **ipx** at the top-level menu, **interface** at the **ipx** menu, and then **define** at the interface menu.

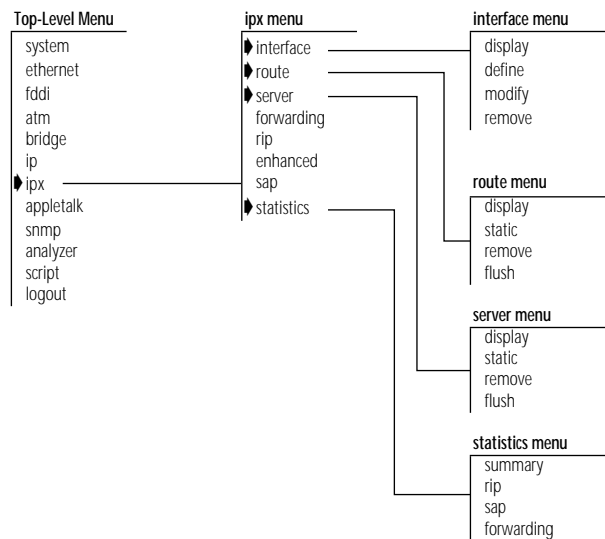


Figure 1-3 IPX Menu Hierarchy

**Appletalk Menu** From the **appletalk** menu, you can view information about and configure Appletalk interfaces, routes, and zones. You can also administer the Appletalk Address Resolution Protocol (AARP), AppleTalk forwarding, and statistics. Figure 1-4 shows the Appletalk menu. For example, to define a new AppleTalk interface, you would enter **appletalk** at the top-level menu, **interface** at the AppleTalk menu, then **define** at the interface menu.

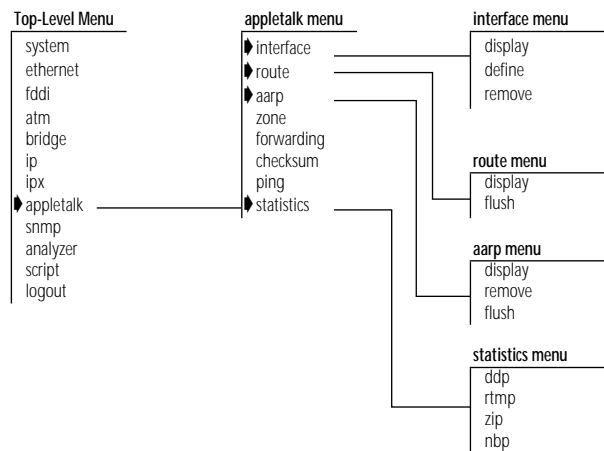


Figure 1-4 Appletalk Menu Hierarchy



# VLANs ON THE LANPLEX® SYSTEM

This chapter contains:

- A description of Virtual LAN (VLAN) concepts and their operational aspects in the LANplex® 2500 system
- Examples of VLAN configurations

---

## About VLANs

The VLAN concept in LAN technology helps minimize broadcast and multicast traffic. It also makes end-station moves, adds, and changes easier for the network administrator.

In the LANplex system, VLANs allow you to:

- Create independent broadcast domains to optimize network performance and create firewalls
- Form flexible user groups independent of the users' physical network location

## Types of VLANs

You can use several types of VLANs to group users. These types include:

- Port group VLANs
- MAC address group VLANs
- Application-oriented VLANs
- Protocol-sensitive VLANs

### Port Group VLANs

Port group VLANs group together one or more switch ports. This simple implementation of VLANs requires little configuration. All frames received on a port are grouped together. For example, all frames received on a port that is part of a port group are kept within that port group, regardless of

the data contained in the frames. Port groups are useful when traffic patterns are known to be directly associated with particular ports. They can benefit the user by restricting traffic based on a set of simple rules.

### MAC Address Group VLANs

VLANs allow a switch to make filtering decisions based on grouping MAC addresses together. These MAC address groups can be configured so that stations in the group can only communicate with each other or with specific network resources. This solution is good for security. It allows the VLAN association to move with the station. However, MAC-address-grouped VLANs may require complex configuration in comparison to other types of VLANs.



*Port group and MAC address group VLANs are supported using the packet filtering capabilities in the LANplex system. For information on port group and MAC address group filtering, refer to your LANplex Operation Guide and LANplex Administration Console User Guide.*

### Application-Oriented VLANs

Using the LANplex filtering capability, application-specific traffic such as telnet traffic or FTP traffic can be filtered based on higher-layer information. You create this application-oriented VLAN by configuring packet filters that specify data and offsets of the data within received packets. For example, to use a filter on a particular port for all telnet traffic, create a filter that discards all TCP traffic received on the telnet port.

IP multicast routing and autocast VLANs are additional VLAN features in the LANplex that can be used to group IP multicast traffic for specific applications. For more information on how the LANplex system manages IP Multicast traffic, see Chapter 8.

### Protocol-Sensitive VLANs

When the LANplex system receives data that has a broadcast, multicast, or unknown destination address, it forwards the data to all ports. This process is referred to as bridge flooding.

Protocol-sensitive VLANs group one or more switch ports together for a specified network layer 3 protocol, such as IP or AppleTalk. These VLANs make flooding decisions based on the network layer protocol of the frame. In addition, for IP VLANs, you can also make flooding decisions based on



layer 3 subnet address information. Protocol-sensitive VLANs allow the restriction of flood traffic for both routable and nonroutable protocols. They have a relatively simple configuration comprising one or more protocols and groups of switch ports. These protocol-sensitive VLANs operate independent of each other. Additionally, the same switch port can belong to multiple VLANs. For example, you can assign port 1 on a LANplex to several IP subnet VLANs, plus one IPX VLAN, one AppleTalk VLAN, and one NetBIOS VLAN. In a multiprotocol environment, protocol-sensitive VLANs can be very effective for controlling broadcast and multicast flooding.



*Two or more types of VLANs can coexist in the LANplex system. When associating received data with a particular VLAN configuration in a multiple VLAN configuration, port group, MAC address group, and application-oriented VLANs always take precedence over protocol-sensitive VLANs.*

### **LANplex Protocol-Sensitive VLAN Configuration**

The LANplex protocol-sensitive VLAN configuration includes three elements: protocol suite, switch ports, layer 3 addressing information for IP VLANs.

#### **Protocol Suite**

The protocol suite describes which protocol entities can comprise a protocol-sensitive VLAN. For example, LANplex VLANs support the IP protocol suite, which is made up of the IP, ARP, and RARP protocols. Table 2-1 lists the protocol suites that the LANplex supports, as well as the protocol types included in each protocol suite.

**Table 2-1** Supported Protocols for VLAN Configuration

| Protocol Suite | Protocol Types   |
|----------------|--|
| IP             | IP, ARP, RARP (Ethernet, SNAP PID)                                       |
| Novell® IPX    | IPX (Ethernet, DSAP, SNAP PID)   |
| AppleTalk®     | DDP, AARP (Ethernet, SNAP PID)   |
| Xerox® XNS     | XNS IDP, XNS Address Translation, XNS Compatibility (Ethernet, SNAP PID) |
| DECnet™        | DEC MOP, DEC Phase IV, DEC LAT, DEC LAVC (Ethernet, SNAP PID)            |
| SNA            | SNA Services over Ethernet (Ethernet)                                    |
| Banyan VINES®  | Banyan (Ethernet, DSAP, SNAP PID)  |

continued

**Table 2-1** Supported Protocols for VLAN Configuration (continued)

| Protocol Suite | Protocol Types               |
|----------------|------------------------------|
| X25            | X.25 Layer 3 (Ethertype)     |
| NetBIOS™       | NetBIOS (DSAP)               |
| Default        | Default (all protocol types) |

### Switch Ports

A group of switch ports is any combination of switch ports on the LANplex system. Included are switch ports created as ATM LAN Emulation Clients (ATM LECs). VLANs do not support media implementations that do not run over switch (bridge) ports, for example, ATM Logical IP Subnets (ATM LISSs).

### Layer 3 Addressing Information

For IP VLANs only, the LANplex system optionally supports configuring of individual IP VLANs with network layer subnet addresses. With this additional layer 3 information, you can create independent IP VLANs that share the same switch ports for multiple IP VLANs. Data is flooded according to both the protocol (IP) and the layer 3 information in the IP header to distinguish among multiple IP VLANs on the same switch port. This configuration is discussed later in the section "Overlapped IP VLANs."

**Default VLAN** When you start up the LANplex system, the system automatically creates a VLAN interface called the default VLAN. Initially, the default VLAN includes all of the switch ports in the system. In the LANplex system, the default VLAN serves to define:

- The flood domain for protocols not supported by any VLAN in the system
- The flood domain for protocols supported by a VLAN in the system but received on nonmember ports

Both cases represent exception flooding conditions that are described in the following sections.

### Modifying the Default VLAN

New switch ports can dynamically appear in the LANplex system if you insert a daughter LAN card or create an ATM LEC. When a new switch port that is not part of a default VLAN appears in the system at initialization, the system software adds that switch port to the first default VLAN defined in the system.



*LANplex VLANs also allow you to modify the initial default VLAN to form two or more subsets of switch ports. If you remove the default VLAN and no other VLANs are defined for the system, no flooding of traffic can occur.*

### How the LANplex® System Makes Flooding Decisions

Protocol-sensitive VLANs directly affect how the LANplex system performs flooding. Without protocol-sensitive VLANs, the flooding process is to forward data to all switch ports in the system. With protocol-sensitive VLANs, the flooding process follows this model:

- As a frame is received that needs to be flooded, it is decoded to determine its protocol type.
- If a VLAN exists for that protocol in the LANplex system and the frame's source port is a member of the VLAN, the frame is flooded according to the group of ports assigned to that VLAN.
- If a VLAN exists for that protocol in the LANplex system but the frame's source port is not a member of the VLAN definition, then the frame is flooded according to the default VLAN assigned to that port.
- If the protocol type of the received frame has no VLAN defined for it in the system, the frame is flooded to the Default VLAN for the receive port.

This example shows how flooding decisions are made according to VLANs set up by protocol (assuming an 18-port switch):

| Index | VLAN    | Ports   |
|-------|---------|---------|
| 1     | Default | 1 - 18  |
| 2     | IP      | 1 - 12  |
| 3     | IPX     | 11 - 16 |

| Data received on... | Is flooded on... | Because...   |
|---------------------|------------------|--|
| IP - port 1         | VLAN 2           | IP data received matches IP VLAN on the source port.                     |
| IPX - port 11       | VLAN 3           | IPX data received matches IPX VLAN on the source port.                   |
| XNS - port 1        | VLAN 1           | XNS data received matches no protocol VLAN, so the Default VLAN is used. |

### VLAN Exception Flooding

If data arrives on a switch port for a certain protocol and VLANs for that protocol are defined in the system but not on that switch port, the default VLAN defines the flooding domain for that data. This case is called VLAN exception flooding.

This example shows how the VLAN exception flooding decision is made (assuming an 18-port switch):

| Index | VLAN    | Ports  |
|-------|---------|--------|
| 1     | Default | 1 - 18 |
| 2     | IP      | 1 - 10 |

| Data received on... | Is flooded on... | Because...  |
|---------------------|------------------|---|
| XNS - port 1        | VLAN 1           | XNS data does not match any defined VLAN in the system.   |
| IP - port 2         | VLAN 2           | IP data received matches IP VLAN 2 for source ports 1 - 10.   |
| IP - port 12        | VLAN 1           | IP data received on source port 12 does not match any defined source port for IP VLAN, so the Default VLAN is used. |

### Overlapped IP VLANs

The LANplex system also gives you the ability to assign network layer information to IP VLANs. This capability allows network administrators to manage their VLANs by subnet. Flooding decisions are made by first matching the incoming frame using the protocol (IP) and then matching it with layer 3 subnet information. If received data is IP but does not match any defined IP subnet VLAN, it is flooded within all IP VLANs using the relevant switch port.

For example, two IP VLANs can be configured for ports 1-10 as follows:

IP VLAN 1 - Subnet 158.101.112.0, ports 1-10

IP VLAN 2 - Subnet 158.101.113.0, ports 1-10

This example shows how flooding decisions are made using overlapping IP VLANs (assuming a 12-port switch):

| Index | VLAN    | Network Address/Mask            | Ports  |
|-------|---------|---------------------------------|--------|
| 1     | Default | none                            | 1 - 12 |
| 2     | IP      | 158.103.122.0/<br>255.255.255.0 | 1 - 6  |
| 3     | IP      | 158.103.123.0/<br>255.255.255.0 | 6 - 12 |

| Data received on...                     | Is flooded on...     | Because...  |
|---|----------------------|---|
| IP subnet<br>158.103.122.2<br>on port 6 | VLAN 2               | IP network layer matches layer 3 address for VLAN 2.              |
| IP subnet<br>158.103.123.2<br>on port 6 | VLAN 3               | IP network layer matches layer 3 address for VLAN 3.              |
| IP subnet<br>158.103.124.2<br>on port 6 | VLAN 2 and<br>VLAN 3 | IP network layer does not match any layer 3 address for IP VLANs. |
| IPX on port 6                           | VLAN 1               | IPX frame does not match any defined VLAN.                        |

As shown in this example, when the subnet address of an IP packet does not match any subnet address of any defined IP VLAN in the system, it is flooded to all of the IP VLANs that share the source switch port, in this case, port 6.

### Routing Between VLANs

The only way for stations that are in two different VLANs to communicate is to route between them. The LANplex system supports internal routing among IP, IPX, and AppleTalk VLANs. If VLANs are configured for other routable network layer protocols, they can communicate between them only via an external router.

The LANplex routing model lets you configure routing protocol interfaces based on a VLAN defined for that protocol. To assign a routing interface, you must first create a VLAN for that protocol and then associate it with that interface.

For example, to create an IP interface that can route through a VLAN:

- 1 Create an IP VLAN for a group of switch ports.  
This IP VLAN does not need to contain layer 3 information unless you want to further restrict flooding according to the layer 3 subnet address.
- 2 Configure an IP interface with a network address, subnet mask, broadcast address, cost, and type (VLAN). Select an IP VLAN to “bind” to that IP interface.

If layer 3 information is provided in the IP VLAN for which you are configuring an IP interface, the subnet portion of both addresses must be compatible.

For example:

IP VLAN subnet 157.103.54.0 with subnet mask of 255.255.255.0

IP host interface address 157.103.54.254 with subnet mask of 255.255.255.0

Layer 2 (bridging) communication is still possible within an IP VLAN (or router interface) for the group of ports within that IP Interface's IP VLAN. IP data destined for a different IP subnet uses the IP routing interface to get to that different subnet, even if the destination subnet is on a shared port.

## VLAN Examples Example 1

Figure 2-1 is an example of a simple configuration that contains three protocol-sensitive VLANs (2 IP and 1 IPX) that share a high-speed FDDI link. The end-stations and servers are on 10Mbps ports with traffic segregated by protocol. They are only aggregated over the high-speed FDDI link. See .




Figure 2-1 Example of a Protocol-Sensitive VLAN Configuration



### Example 2

Figure 2-2 is an example of a configuration that contains two different protocol-sensitive VLANs (IP and IPX) with servers on separate high-speed 100BASE-T ports. The end-station clients share the same switch ports, yet the IP and IPX traffic stays separate. See Figure 2-2.




Figure 2-2 A VLAN Configuration with Servers on Separate 100BASE-T ports.



# 3

## BRIDGING AND ROUTING IN THE LANPLEX® SYSTEM

This chapter shows how the LANplex® system operates in a subnetworked routing environment and describes the LANplex routing methodology — specifically, how the LANplex bridging and routing model compares with traditional models.

---

### What Is Routing?

Routing is the process of distributing packets over potentially dissimilar networks. A router (also called a gateway) is the machine that accomplishes this task. Routers are typically used to:

- Connect enterprise networks together
- Connect subnetworks (or client/server networks) to the enterprise network

Figure 3-1 shows where routers are typically used in a network.

The LANplex system performs routing that connects subnets to the enterprise network, providing connectivity between devices within a workgroup, department, or building.




Figure 3-1 Traditional Architecture of a Routed Network

### LANplex in a Subnetworked Environment

The LANplex system allows you to fit Ethernet switching capability into highly subnetworked environments. When you put the LANplex system into such a network, the system streamlines your network architecture and easily switches traffic between and within subnets over Ethernet and FDDI. See Figure 3-2.





Figure 3-2 Subnetted Architecture with LANplex® Switching Hubs

### Integrating Bridging and Routing

The LANplex system integrates bridging and routing. Multiple switch ports can be assigned to each subnet. See Figure 3-3. Traffic between ports assigned to the same subnet is switched transparently using transparent bridging or Express switching (described in the *LANplex® 2500 Operation Guide*). Traffic traveling to different subnets is routed using one of the supported routing protocols.



*In the following descriptions of bridging and routing on the LANplex system, the term **MAC address** refers to a physical hardware address. The term **network address** refers to a logical address that applies to a specific protocol.*



**Figure 3-3** Multiple Ports per Subnets with the LANplex 2500 System

Because the LANplex model of bridging and routing allows several segments to be connected to the same subnet, you can increase the level of segmentation in your network without having to create new subnets or assign network addresses. Instead, you can use additional Ethernet ports to expand your existing subnets. This is in contrast to more traditional forms of bridging and routing where, at most, one port is connected to any subnet.

In the traditional model, if you want to increase the level of segmentation in your network, you must create additional subnets and assign new network addresses to your existing hosts.

---

## Bridging and Routing Models

The way routing is implemented in the LANplex system differs from how bridging and routing usually coexist in a system.

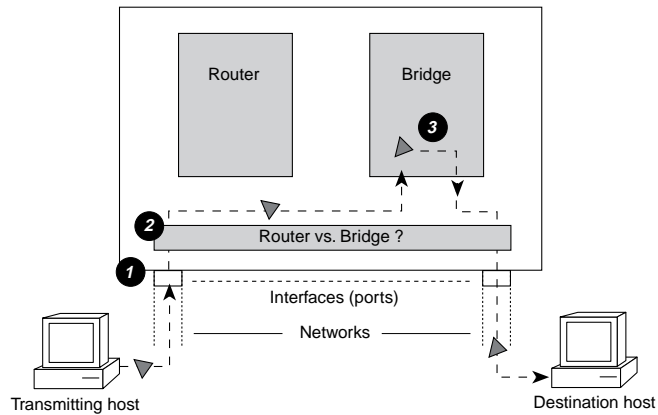
- **Traditional Bridging and Routing Model** — Traditionally, bridging and routing are peer entities; either a packet is bridged or routed. Packets belonging to recognized protocols are routed; all others are bridged.
- **LANplex Bridging and Routing Model** — In the LANplex model, the bridge and router operate hierarchically on the LANplex system, routing over bridging. When a packet enters the system, the system first tries to bridge the packet. If the packet's destination network address is not on the same subnet, then the system routes the packet.

## Traditional Bridging and Routing Model

The bridge or router determines whether a packet should be bridged or routed based on the protocol to which the packet belongs. If the packet belongs to a recognized protocol, the packet is routed. Otherwise, it is bridged.

In the traditional bridging and routing model, a packet is *bridged* as follows (see Figure 3-4):


- 1 The packet enters the bridge or router.
- 2 The bridge or router determines that the packet does *not* belong to a recognized routing protocol, so the packet is passed to the bridge.
- 3 The bridge examines the destination MAC address and forwards the packet to the port on which that address has been learned.



**Figure 3-4** Bridging in the Traditional Bridging and Routing Model

In the traditional bridging and routing model, a packet is *routed* as follows (see Figure 3-5):

- 1 The packet enters the bridge or router.
- 2 The bridge or router determines that the packet belongs to a recognized routing protocol, so the packet is passed to the router.
- 3 The router examines the destination network address and forwards the packet to the interface (port) connected to the destination subnet.



**Figure 3-5** Routing in the Traditional Bridging and Routing Model

### LANplex Bridging and Routing Model

The LANplex 2500 system uses the destination MAC address to determine whether it will bridge or route a packet. Before a host system sends a packet to another host, it compares its own network address to the network address of the other host as follows:

- If network addresses are on the same subnet, the packet is bridged directly to the destination host's address.
- If network addresses are on different subnets, the packet must be routed from one subnet to the other. In this case, the host transmits the packet to the connecting router's MAC address.

In the LANplex bridging/routing model, a packet is *bridged* as follows (see Figure 3-6):

- 1 The packet enters the LANplex system.
- 2 The packet's destination MAC address is examined by the bridging layer.
- 3 The destination MAC address does not correspond to the MAC address of one of the system ports configured for routing. The bridging layer selects a segment (port) based on the destination MAC address and forwards the packet to that segment.




Figure 3-6 Bridging in the LANplex Bridging and Routing Model



In the LANplex bridging and routing model, a packet is *routed* as follows (see Figure 3-7):

- 1 The packet enters the LANplex system.
- 2 The packet's destination address is examined by the bridging layer.
- 3 The destination address corresponds to the address of one of the system ports configured for routing (as opposed to a learned end-station address). The packet is passed to the router interface associated with the port on which the packet was received.
- 4 The routing layer:
  - a Selects a destination interface based on the destination network address.
  - b Determines the MAC address of the next hop (either the destination host or another gateway).
  - c Passes the packet back to the bridging layer.
- 5 The bridging layer then selects a segment (port) based on the destination MAC address and forwards the packet to that segment.




Figure 3-7 Routing in the LANplex Bridging and Routing Model



# 4

## ROUTING WITH IP TECHNOLOGY

This chapter gives an overview of IP routing technology, specifically defining:

- What IP routing involves
- What elements are necessary for IP routers to effectively transmit packets
- How IP routing transmission errors are detected and resolved
- Routing with classical IP over ATM

### IP Routing and the OSI Model

An IP router, unlike a bridge, operates at the network layer of the OSI Reference Model. That is, it routes packets by examining the network layer address (IP address). Bridges use the data-link layer MAC addresses to make forwarding decisions. See Figure 4-1.

#### OSI Reference Model

Application Layer

Presentation Layer

Session Layer

Transport Layer

Network Layer

IP

RIP

ARP

ICMP

Data-link Layer

MAC

Physical Layer

Figure 4-1 OSI Reference Model and IP Routing

When an IP router sends a packet, it does not know the complete path to a destination — only the next hop. Each hop involves three steps:

- The IP routing algorithm computes the *next hop* IP address, and next router interface, using the routing table entries.
- The Address Resolution Protocol (ARP) translates the next hop IP address into a physical MAC address.
- The router sends the packet over the network to the next hop.

These routing elements are described in more detail in the following section.

---


## Elements of IP Routing

IP routers use the following elements to transmit packets in a subnetworking environment:

- IP addresses
- Router interfaces
- Routing tables
- Address Resolution Protocol (ARP)

## IP Addresses

IP addresses are 32-bit addresses composed of a *network part* (the address of the network on which the host is located) and a *host part* (the address of the host on that network). See Figure 4-2. IP addresses differ from Ethernet and FDDI MAC addresses, which are unique hardware-configured 48-bit addresses.



**Figure 4-2** IP Address: Network Part and Host Part

A central agency assigns the network part of the IP address, and the network administrator assigns the host part. All devices connected to the same network share the same IP address prefix (the network part of the address).

## Address Classes

The boundary of the network part and the host part depends on the class that the central agency assigns to your network. The primary classes of IP addresses are Class A, Class B, and Class C.

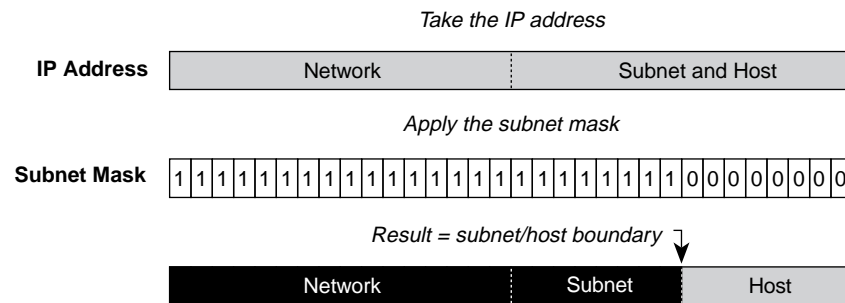
- **Class A addresses** — have 8 bits for the network part and 24 bits for the host part. Although only a few Class A networks can be created, each can contain a very large number of hosts.
- **Class B addresses** — have 16 bits for the network part and 16 bits for the host part.
- **Class C addresses** — have 24 bits for the network part and eight bits for the host part. Each Class C network can contain only up to 254 hosts, but many such networks can be created.

The class of an IP address is designated in the high-order bits of the network parts of the address.

## Subnet Part of an IP Address

In some environments, the IP address contains a *subnet part*. Subnetting allows a single Class A, B, or C network to be further subdivided internally while still appearing as a single network to other networks. The subnet part of the IP address is only visible to those hosts and gateways on the subnet network.

When an IP address contains a subnet part, a *subnet mask* is used to identify which bits are the subnet address and which are the host address. A subnet mask is a 32-bit number that uses the same format and representation as IP addresses. Each IP address bit corresponding to a *1* in the subnet mask is in the network or subnet part of the address. Each IP address bit corresponding to a *0* is in the host part of the IP address. See Figure 4-3.



**Figure 4-3** How a Subnet Mask Is Applied to the IP Address

An example of an IP address that includes network, subnet, and host parts is *158.101.230.52* with a subnet mask of *255.255.255.0*. This address is divided as follows:

- *158.101* is the network part
- *230* is the subnet part
- *52* is the host part

### Router Interfaces

A router interface is the connection between the router and a subnet. In traditional routing models, the interface is the same as the port, since only one interface can exist per port. In the LANplex system's IP routing model, more than one port can be connected to the same subnet.

Each router interface has an IP address and a subnet mask. This address defines both the number of the network to which the router interface is attached and its host number on that network. A router interface's IP address serves two functions:

- The IP address is used when sending IP packets to or from the router itself.
- The IP address defines the network and subnet numbers of the segment connected to that interface. See Figure 4-4.




Figure 4-4 Router Interfaces in the LANplex System

## Routing Table

A routing table allows a router or host to determine how to send a packet toward the packet's ultimate destination. The routing table contains an entry for every destination network, subnet, or host to which the router or host is capable of forwarding packets. A router or host uses the routing table when the destination IP address of the packet it is sending is not on a network or subnet to which it is directly connected. The routing table provides the IP address of a router that can forward the packet toward its destination.

The routing table consists of the following elements:

- **Destination IP Address** — the destination network, subnet, or host
- **Subnet Mask** — the subnet mask corresponding to the destination IP address
- **Metric** — a measure of the “distance” to the destination. In the Routing Information Protocol (RIP), the metric is the number of hops.
- **Gateway** — the IP address of the next hop router (the IP address of the interface through which the packet travels)
- **Interface** — the interface number through which a packet must travel to reach that router

Figure 4-5 shows the routing table of the router in Figure 4-4.

| Routing Table          |               |        |             |           |
|------------------------|---------------|--------|-------------|-----------|
| Destination IP Address | Subnet Mask   | Metric | Gateway     | Interface |
| 158.101.1.1            | 255.255.255.0 | 1      | 158.101.1.2 | 1         |
| 158.101.2.1            | 255.255.255.0 | 1      | 158.101.2.2 | 2         |
| 158.101.3.1            | 255.255.255.0 | 1      | 158.101.3.2 | 3         |
| default route          | 255.255.255.0 | 1      | 158.101.1.2 | 1         |

Figure 4-5 Example of a Routing Table in the LANplex Routing Model

Routing table information is generated and updated in either of the following ways:

- **Statically** — You manually enter routes, which do not change until you change them (that is, they will not time out).
- **Dynamically** — The router uses a routing protocol, such as RIP, to exchange information. Routes are recalculated at regular intervals.

### Static Routes

A static route is one that you manually configure in the routing table. Static routes are useful in environments where no routing protocol is used or where you want to override some of the routes generated with a routing protocol. Because static routes do not automatically change in response to network topology changes, you should manually configure only a small number of reasonably stable routes.

### Dynamic Routes Using RIP

Automated methods of configuring routes help you keep up with a changing network environment, allowing routes to be reconfigured quickly and reliably. Interior Gateway Protocols (IGP), which operate within networks, provide this automated method. The LANplex system uses the Routing Information Protocol (RIP), one of the most widely used IGPs, to configure its routing tables dynamically.

RIP operates in terms of active and passive devices. The *active devices*, usually routers, broadcast their RIP messages to all devices in a network or subnet; they update their own routing tables when they receive a RIP message. The *passive devices*, usually hosts, listen for RIP messages and update their routing tables; they do not send RIP messages.



An active router sends a RIP message every 30 seconds. This message contains both the IP address and a metric (the distance to the destination from that router) for each destination. In RIP, each router that a packet must travel through to reach a destination equals one *hop*.

### Default Route

In addition to the routes to specific destinations, the routing table may contain an entry called the *default route*. The router uses the default route to forward packets that do not match any other routing table entry. A default route is often used in place of routes to numerous destinations all having the same gateway IP address and interface number. The default route can be configured statically, or it can be learned dynamically using RIP.

### Address Resolution Protocol (ARP)

ARP is a low-level protocol used to locate the MAC address corresponding to a given IP address. This protocol allows a host or router to make its routing decisions using IP addresses while it uses MAC addresses to forward packets from one hop to the next.

Once the host or router knows the IP address of the *next hop* to the destination, the host or router must translate that IP address into a MAC address before the packet can be sent. To do this translation, the host or router first looks in its ARP cache, a table of IP addresses with their corresponding MAC addresses. Each device participating in IP routing maintains an ARP cache. See Figure 4-6.

| ARP Cache   |              |
|-------------|--------------|
| IP Address  | MAC Address  |
| 158.101.1.1 | 00308e3d0042 |
| 158.101.2.1 | 0080232b00ab |

**Figure 4-6** Example of an ARP Cache

If the IP address does not have a corresponding MAC address listed, the host or router broadcasts an *ARP request* packet to all the devices on the network. The ARP request contains information about the hardware and

protocol. The two key elements of the ARP request are the target and source addresses for both the hardware (MAC addresses) and the protocol (IP addresses). See Figure 4-7.

#### ARP Request

|               |                         |
|---------------|-------------------------|
|               |                         |
| 00802322b00ad | Source hardware address |
| 158.101.2.1   | Source protocol address |
| ?             | Target hardware address |
| 158.101.2.15  | Target protocol address |

**Figure 4-7** Example of an ARP Request Packet

When the devices on the network receive this packet, they examine it, and if their address is not the target protocol address, they discard the packet. When a device receives the packet and confirms that its IP address matches the target protocol address, this device places its MAC address in the target hardware address field and sends the packet back to the source hardware address. When the originating host or router receives the *ARP reply*, it places the new MAC address in its ARP cache next to the corresponding IP address. See Figure 4-8.

#### ARP Cache

| IP Address  | MAC Address  |
|-------------|--------------|
| 158.101.1.1 | 00308e3d0042 |
| 158.101.2.1 | 0080232b00ab |
| 158.101.3.1 | 0134650f3000 |

**Figure 4-8** Example of ARP Cache Updated with ARP Reply

Once the MAC address is known, the host or router can send the packet directly to the next hop.

---

## IP Routing Transmission Errors

Because each router only knows about the next hop, it is not aware of problems that might be further “down the road” toward the destination. Destinations can be unreachable if:

- Hardware is temporarily out of service
- You inadvertently specified a nonexistent destination address
- The router does not have a route to the destination network

To help routers and hosts know of problems in packet transmission, an error-reporting mechanism called Internet Control Message Protocol (ICMP) provides error reporting back to the source when routing problems arise. ICMP allows you to determine whether a delivery failure resulted from a local or a remote malfunction.

ICMP does the following:

- Tests the reachability of nodes (*ICMP Echo Request* and *ICMP Echo Reply*)

A host or gateway sends an ICMP echo request to a specified destination. If the destination receives the echo request, it sends an ICMP echo reply back to the original sender. This process tests whether the destination is reachable and responding and verifies that the major pieces of the transport system work. The *ping* command is one frequently used way to invoke this process.
- Creates more efficient routing (*ICMP Redirect*)

Often the host route configuration specifies the minimal possible routing information needed to communicate (for example, the address of a single router). The host relies on routers to update its routing table. In the process of routing packets, a router may detect a host not using the best route. The router then sends the host an ICMP redirect, requesting that the host use a different gateway when sending packets to that destination. The next time the host sends a packet to that same destination, it uses the new route.
- Informs sources that a packet has exceeded its allocated time to exist within the network (*ICMP Time Exceeded*)

---

## Routing with Classical IP over ATM

LANplex Extended Switching software supports classical IP routing over ATM ARP in an ATM network. Classical IP over ATM uses Logical IP Subnets (LISs) to forward packets within the network environment.



*See the LANplex® 2500 Operation Guide for detailed information about the ATM protocol architecture. See the LANplex® 2500 Administration Console User Guide for information about how to configure ATM ports.*

### About Logical IP Subnets (LISs)

An LIS is a group of IP nodes that belong to the same subnet, and which are directly connected to a single ATM network. When you add a node to a LIS through the Administration Console IP interface menu, you define its IP address, subnet mask, and the address an ATM ARP server that supports it.

### ATM ARP Servers

An ATM ARP server maintains a table of IP addresses and their corresponding ATM addresses and circuit information. To forward IP packets over an ATM interface, the network node learns the ATM address for the corresponding IP address from the ATM ARP server.

Each ATM ARP server supports a single LIS. You can associated two or more LISs with the same ATM network, but each LIS operates independently of other LISs on the network.

Several types of network nodes can function as ATM ARP servers:

- Any LANplex system with revision 8.1.0 or later of Extended Switching software
- An ATM switch
- A UNIX® workstation

The following sequence describes how the ATM ARP server learns and stores information about the IP and ATM addresses of nodes in the network.

- A node establishes a connection to the ATM ARP server
- The ATM ARP server sends an inverse ATM ARP request to the node, requesting its IP and ATM address
- When the node returns this information, the ATM ARP server stores, or *caches*, it in the ATM ARP server table.

### Forwarding to Nodes within an LIS

Nodes can forward packets directly to other nodes in the same LIS. To forward a packet within the same LIS, the sending node requests a translation from the destination IP address to the corresponding ATM address from the ATM ARP server.

- If the address is known to the server, the server returns a message with this address
- If the address is not known to the server, the server returns a message to advise the sending node that the packet is discarded.

When the server returns a destination address, the sending node uses this learned address to create a virtual circuit (VC) and to forward this and all subsequent packets to the destination address. The sending node adds this VC to its ATM ARP cache.

---

## IP Routing References

Comer, Douglas E. *Internetworking with TCP/IP. Volume I: Principles, Protocols, and Architecture*. Englewood Cliffs, New Jersey: Prentice Hall, Inc., 1991.

Perlman, Radia. *Interconnections: Bridges and Routers*. Reading, Massachusetts: Addison-Wesley Publishing Company, Inc., 1992.

Sterns, Richard. *TCP/IP Illustrated. Volume 1: The Protocols*. Reading, Massachusetts: Addison-Wesley Professional Computing Services, 1992.

RFC 791. *Internet Protocol Specification*.

RFC 792. *Internet Control Message Protocol Specification*.

RFC 1009. *Requirements for Internet Gateways*.

RFC 1042. *A Standard for the Transmission of IP Datagrams over IEEE 802 Networks*.

RFC 1058. *Routing Information Protocol*.

RFC 1122. *Requirements for Internet Hosts*.

RFC 1577. *Classical IP over ATM*.



# 5

## ROUTING WITH IP MULTICAST

This chapter describes the IP multicast routing implementation on the LANplex® system.

---

### About IP Multicast Routing

IP multicast routing is an extension of the Internet Protocol. Multicast routing allows a router or switch to send packets to a specific group of hosts without using broadcasts or multiple unicast transmissions. This group can include members that reside on the local LAN, members that reside on different sites within a private network, or members that are scattered throughout the Internet. Multicast routing achieves this functionality without loops or excess transmissions.

IP Multicast support within the LANplex system has two main components:

- Internet Group Management Protocol (IGMP)
- Distance Vector Multicast Routing Protocol (DVMRP)

This chapter describes these two protocols as well as the algorithms that the LANplex system uses for multicast routing.

---

### IGMP

The LANplex system is capable of dynamic multicast filtering based on the Internet Group Management Protocol (IGMP). This protocol ensures that multicast packets are flooded only to the appropriate ports within a routing interface.

IGMP tracks end-station group membership within a multicast group. Membership in a group is dynamic, and hosts are allowed to be a member of more than one group at a time. Broadcast domains are maintained by avoiding propagation of multicast broadcasts to the entire subnet by confining them within the group (IGMP "snooping").

---

## DVMRP

The Distance Vector Multicast Routing Protocol (DVMRP) establishes the multicast delivery path over a series of routing devices. DVMRP is a simple distance vector routing protocol, similar to the IP Routing Information Protocol (RIP). Multicast routers exchange distance vector updates that contain lists of destinations as well as the distance in hops to each destination. They maintain this information in a routing table.

DVMRP is the current routing protocol used on the Internet Multicast Backbone (MBONE). Full support of DVMRP allows the LANplex system to fully establish the delivery path without requiring a direct connection to a multicast router.

## The MBONE

The MBONE is an experimental “Multicast Backbone” network that exists on the Internet. Users can test multicast applications and technology on the MBONE without waiting for Internet multicast standards to be set. You can gain access to the MBONE through any Internet service provider.

The MBONE routers forward mulitcast packets over an interface or over a multicast tunnel only if the Time-To-Live (TTL) value present in the packet is larger than the tunnel's threshold. (See the section “Multicast Tunnels” on page 6 for more information about tunnels.)



*LANplex 2500 systems at revisions earlier than 8.0 support up to 16 IP multicast tunnels or routing interfaces when connected to the MBONE network. LANplex 2500 systems at revision 8.0 or later can support up to 32 IP multicast tunnels or routing interfaces when connected to the MBONE.*



---

**Multicast Routing Algorithms**

The LANplex system uses three algorithms that support multicast routing:

- Flooding
- Spanning Trees
- Reverse Path Forwarding

**Flooding**

Several types of flooding algorithms exist, but they all share the same general principles: a node in the network receives a packet that was sent to a multicast destination. The node determines whether the packet is an original that it has not seen before or a duplicate of a packet that it has seen before. If the packet is an original, the node forwards the packet on all interfaces except the incoming interface. If the packet is a duplicate, the node discards it.


The flooding algorithm is useful in situations where the most important requirement for the network is robustness. It does not depend on any kind of routing tables. Destinations will receive packets as long as at least one path to them exists and no errors occur during transmission.

**Spanning Trees**

The Spanning Tree algorithm detects loops and logically blocks redundant paths in the network. The paths form a loopless graph, or tree, spanning all the nodes in the network. A port in the blocking state does not forward or receive data packets.


After the algorithm eliminates extra paths, the network configuration stabilizes. When one or more of the paths in the stable topology fail, the protocol automatically recognizes the changed configuration and activates redundant links. This strategy ensures that all nodes remain connected.

Figure 5-1 shows a simple network with five links.



**Figure 5-1** Simple Network Implemented Without Using Spanning Tree

A spanning tree for this network consists of links 1, 2, 3, and 4. See Figure 5-2.



**Figure 5-2** Spanning Tree Algorithm Implemented to Block Redundant Paths

### Reverse Path Forwarding

Reverse path forwarding (RPF) is the multicast algorithm in use on the MBONE network. RPF is designed to avoid duplicate paths on multi-access links. It uses a routing table to compute a logical spanning tree for each network source. The RPF algorithm has these basic steps:

- 1 When the system receives a multicast packet, the algorithm notes the source network of the packet and the interface on the LANplex system that received the packet.
- 2 If the interface belongs to the shortest path towards the source network, then the system forwards the packet to all interfaces except the interface on which the packet was received.
- 3 If the condition in Step 2 is false, the system drops the packet.

**Pruning** *Pruning* is a method used in the RPF algorithm to forward packets to a spanning tree only if group members exist in the tree. This method results in fewer spanning trees, but it requires dynamic updates to the routing table.

Nodes that are at the border of the network and have no point beyond them in the RPF spanning tree are called *leaf* nodes. Leaf nodes all receive the first multicast packet. If a group member is attached to the leaf node, the node continues to accept packets. If no group member is attached to the leaf node, the node sends back a "prune message" to the router that sent the packet. The message tells the router to send no further packets to this group. In the LANplex system, the Administration Console IP multicast CacheDisplay includes information about when pruning will occur on the spanning tree.

---

**Multicast Interfaces**

Multicast interfaces on the LANplex system have several characteristics which are described in this section:

**DVMRP Metric Value**

The DVMRP metric value determines the cost of a multicast interface. The higher the cost, the slower the link. The default value is 1.

**Time-To-Live (TTL) Threshold**


This TTL threshold determines whether the interface will forward multicast packets to other switches and routers in the subnet. If the interface TTL is greater than the packet TTL, then the interface does not forward the packet. The default value is one 1, which means that the interface forwards most packets.

**Rate Limit** The rate limit determines how many multicast packets can travel over the interface in kilobytes-per-second. The LANplex system drops multicast traffic that travels faster than this rate. The default is set to 0, which implies no rate limit is set. In all other instances, the lower the rate limit, the more limited the traffic over the interface.

---

**Multicast Tunnels** Multicast tunnels are logical connections between two multicast routers through one or more unicast routers. The multicast router at the local endpoint of the tunnel encapsulates multicast packets in a format that unicast routers can interpret and forward. The multicast router at the remote endpoint decapsulates the packets into their multicast format. Tunnels are virtual links through the unicast IP network.

Multicast tunnels have characteristics similar to those of a multicast interface: a DVMRP metric value, a TTL threshold, and a rate limit. When you define a multicast tunnel, you also specify the destination address of the remote multicast router that is the remote endpoint of the tunnel.



## ROUTING WITH IPX

This chapter provides an overview of IPX routing, including:

- What part IPX plays in the NetWare environment
- How IPX works
- What elements are necessary for IPX routers to transmit packets effectively

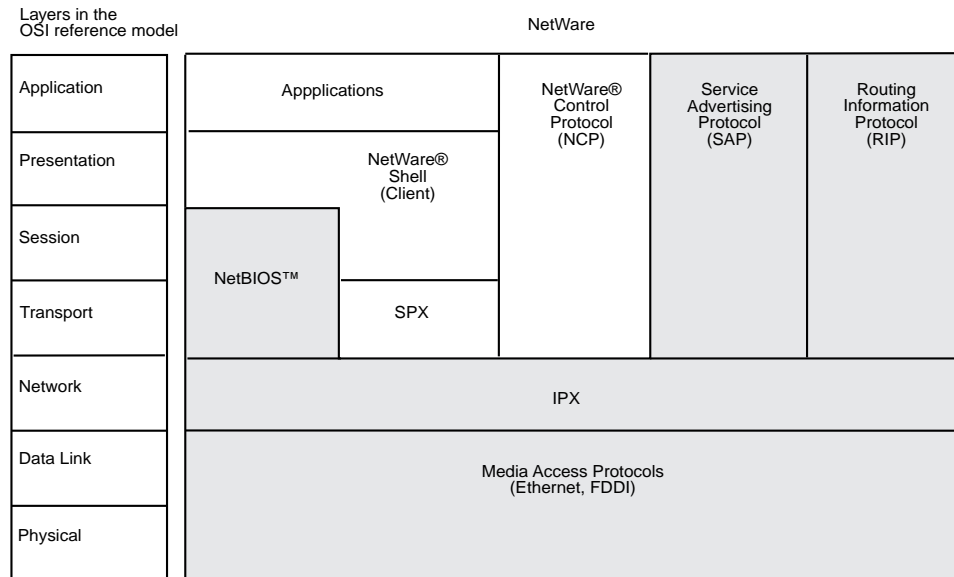
---

### **IPX Routing in the NetWare® Environment**

The NetWare® network operating system was developed and introduced to the market by Novell, Inc. in the early 1980s. Much of the NetWare networking technology was derived from Xerox Network System (XNS)<sup>™</sup>, a networking system developed by Xerox Corporation.

The NetWare operating system is based on a client/server architecture where clients request certain services from servers such as file access and printer access. As a network operating system environment, the NetWare operating system specifies the upper five layers of the OSI reference model. It provides file and printer sharing and supports various applications such as electronic mail and database access.

Figure 6-1 illustrates a simplified view of NetWare's better-known protocols and their relationship to the OSI reference model.



**Figure 6-1** NetWare Protocols and the OSI Reference Model

The LANplex system uses the following protocols for routing in a Netware environment:

- Internet Packet Exchange (IPX)
- Routing Information Protocol (RIP)
- Service Advertisement Protocol (SAP)

### Internet Packet Exchange (IPX)

IPX is the primary protocol used for routing in a netware environment. This datagram, connectionless protocol does not require an acknowledgment for each packet sent. Any packet acknowledgment, or connection control, must be provided by protocols above IPX.

IPX defines internetwork and intranode addressing schemes. IPX internetwork addressing is based on network numbers that are assigned to each interface in an IPX network. IPX intranode addressing is in the form of socket numbers. Since several processes are normally operating within a node, socket numbers provide a type of mail slot so that each process can distinguish itself to IPX.

**Routing  
Information  
Protocol (RIP)**

RIP allows the exchange of routing information on a NetWare network. IPX routers use RIP to dynamically create and maintain their routing tables.

RIP allows one router to exchange routing information with a neighboring router. As a router becomes aware of any changes in the network layout, it broadcasts this information to any neighboring routers. IPX routers also send periodic RIP broadcast packets containing all routing information known to the router. These broadcasts synchronize all routers on the network and age those networks that might become inaccessible if a router becomes disconnected from the network abnormally.

**Service Advertising  
Protocol (SAP)**

SAP provides routers and servers that contain SAP agents with a means of exchanging network service information.

Through SAP, servers advertise their services and addresses. Routers gather this information and share it with other routers. This strategy allows routers to dynamically create and maintain a database (server table) of network service information. Clients on the network can determine what services are available and obtain the network address of the nodes (servers) where they can access those services. Clients require this information to initiate a session with a file server.

SAP allows one router to exchange information with a neighboring SAP agent. As a router's SAP agent becomes aware of any change in the network server layout, it immediately broadcasts this information to any neighboring SAP agents. The router also periodically sends SAP broadcast packets containing all server information known to the SAP agent. These broadcasts synchronize all servers on the network and age those servers that might become inaccessible because of any abnormal shut down of the router or server.

## How IPX Routing Works

A router operates at the network layer of the OSI Reference Model. This means that it receives its instructions to route packets from one segment to another from a network-layer protocol. IPX, with the help of RIP, performs these network layer tasks. These tasks include addressing, routing, and switching information packets to move single packets from one location to another. This section first describes the information included in an IPX packet that helps it get delivered and then it describes the IPX packet delivery process.

### IPX Packet Format

The IPX packet format consists of two parts: a 30-byte header and a data portion. The network, node, and socket address for both the destination and source are held within the packet's IPX header.

Figure 6-2 shows the IPX packet format.

|                               |                      |
|-------------------------------|----------------------|
| Checksum (2 bytes)            |                      |
| Packet Length (2 bytes)       |                      |
| Transport Control (1 byte)    | Packet Type (1 byte) |
| Destination Network (4 bytes) |                      |
| Destination Node (6 bytes)    |                      |
| Destination Socket (2 bytes)  |                      |
| Source Network (4 bytes)      |                      |
| Source Node (6 bytes)         |                      |
| Source Socket (2 bytes)       |                      |
| Upper-layer Data              |                      |

**Figure 6-2** IPX Packet Format



The packet format consists of the following elements:

- **Checksum** — The IPX packet begins with a 16-bit checksum field that is set to 1s.
- **Packet Length** — This 16-bit field contains the length, in bytes, of the complete network packet. This field includes both the IPX header and the data. The IPX length must be at least 30 bytes.
- **Transport Control** — This 1-byte field indicates how many routers a packet has passed through on its way to its destination. Packets are discarded when this value reaches 16. A sending node always sets this field to 0 when building an IPX packet.
- **Packet Type** — This 1-byte field specifies the upper-layer protocol that will receive the packet's information.
- **Destination Network** — This 4-byte field provides the destination node's network number. When a sending node sets this field to zero, the destination node is assumed to be on the same local segment as the sending node.
- **Destination Node** — This 6-byte field contains the physical address of the destination node.
- **Destination Socket** — This 2-byte field contains the socket address of the packet's destination process.
- **Source Network** — This 4-byte field provides the source node's network number. If a sending node sets this field to 0, it means the source's local network is unknown.
- **Source Node** — This 6-byte field contains the physical address of the source node. Broadcast addresses are not allowed.
- **Source Socket** — This 2-byte field contains the socket address of the process that transmitted the packet.
- **Upper-layer Data** — The data field contains information for the upper-layer processes.

**IPX Packet Delivery** On a NetWare network, the successful delivery of a packet depends both on the proper addressing of the packet and on the internetwork configuration. Packet addressing is handled in the packet's Media Access Control (MAC) protocol header and IPX header address fields.

To send a packet to another node, the sending node must know the complete internetwork address including the network, node, and socket of the destination node. Once the sending node has the destination node's address, it can proceed with addressing the packet. However, the way the MAC header of that packet is addressed depends on whether the sending and destination nodes are separated by a router. See Figure 6-3.

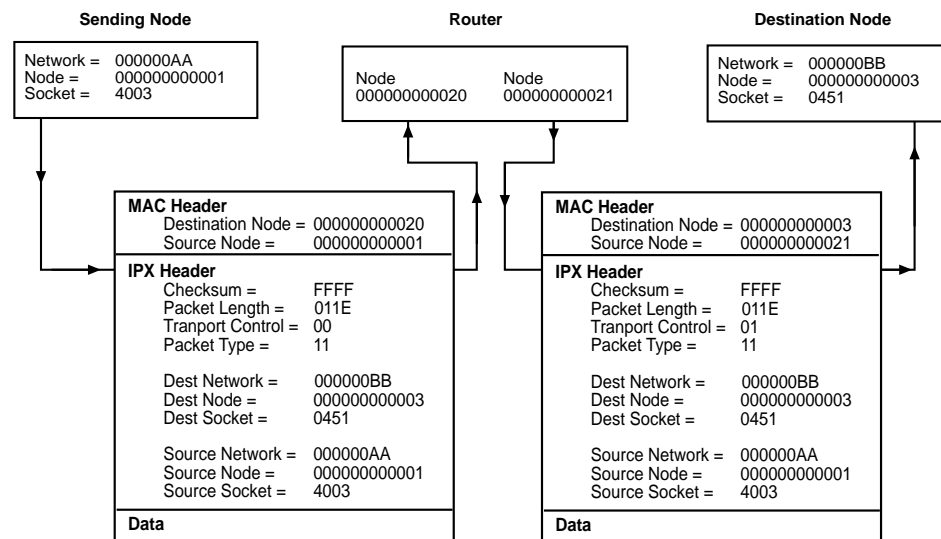


Figure 6-3 IPX Packet Routing

### Sending Node's Responsibility

When a node needs to send information to another node with the same network number, the sending node can simply address and send packets directly to the destination node. However, if the sending and receiving nodes have different network numbers, the sending node must find a router on its own network segment that can forward packets to the destination node's network segment.

To find this router, the sending node broadcasts a RIP packet requesting the best route to the destination node's network number. The router residing on the sending node's segment with the shortest path to the destination segment responds to the RIP request. The router's response includes its network and node address in the IPX header. If the sending node is a router rather than a workstation, the router can get this information from its internal routing tables and need not send a RIP request.

Once the sending node knows the router's node address, it can send packets to the destination node.

### **Router's Responsibility**

When a router receives an IPX packet, it handles the packet in one of two ways:

- If the packet is destined for a network number to which the router is directly connected, the router performs the following tasks:
  - Places the destination node address from the IPX header in the destination address field of the MAC header.
  - Places its own node address in the source address field of the MAC header.
  - Increments the Transport Control field of the IPX header and transmits the packet on the destination node segment.
- If the packet is destined for a network number to which the router is not directly connected, the router sends the packet to the next router along the path to the destination node as follows:
  - The router looks up the node address (in the routing information table) of the next router and places the address in the destination address field of the packet's MAC header. For more information on routing tables, see the next section.
  - The router places its own node address in the source address field of the packet's MAC header.
  - The router increments the Transport Control field in the IPX header and sends the packet to the next router.

---

## The Elements of IPX Routing

IPX routers use the following elements to transmit packets over an intranetwork:

- Router interfaces
- Routing tables
- Service Advertising Protocol (SAP)

### Router Interfaces

A router interface is the connection between the router and the network number (address). In traditional routing models, the interface would be the same as the port, because only one interface can exist per port.

In the LANplex system's IPX routing, more than one port can be connected to the network number. Therefore, the router interface is the relationship between the ports and the network number (address) in your IPX network.

Each router interface has a network address. This address defines the network number to which the router interface is attached. The router interface's IPX address serves two functions:

- It is used when sending IPX packets to or from the router itself.
- It defines the network number of the segment connected to that interface.

### Routing Tables

A routing table holds information about all the network segments. It allows a router to send a packet toward its ultimate destination using the best possible route. The routing information table contains an entry for every network number that the router currently knows exists. A router uses the routing information table when the destination network number of the packet it is sending is not on a network to which it is directly connected. The routing information table provides the immediate address of a forwarding router that *can* forward the packet toward its destination.

The routing table consists of the following elements:

- **Interface** — Identifies the number of the router's interface that will be used to reach the specific network segment.
- **Address** — Identifies the addresses for segments that the router currently knows exists.

- **Hops to Network** — Provides the number of routers that must be crossed to reach the network segment.
- **Ticks to Network** — Provides an estimate of the time necessary to reach the destination segment.
- **Node** — The node address of the router that can forward packets to each segment. When set to all zeroes, the route is directly connected.
- **Aging Timer** — The time since the network's last update.

Figure 6-4 shows an example of a typical routing information table.

| Routing Table |          |      |       |                   |     |
|---------------|----------|------|-------|-------------------|-----|
| Interface     | Address  | Hops | Ticks | Node              | Age |
| 1             | 1        | 1    | 1     | 00-00-00-00-00-00 | 0   |
| 2             | 45469f30 | 1    | 1     | 00-00-00-00-00-00 | 0   |
| 2             | 45469f33 | 2    | 3     | 08-00-17-04-33-45 | 40  |

Figure 6-4 Routing Table Example

### Generating Routing Table Information

The routing information table is generated and updated as follows:

- **Statically** — You manually enter routes. They do not change until you change them (they do not time out).
- **Dynamically** — The router uses RIP to exchange information with other routers. Routes are recalculated at regular intervals.

**Static Routes.** A static route is one you manually configure in the routing table. Static routes are useful in environments where no routing protocol is used or where you want to override some of the routes generated with a routing protocol. Because static routes do not automatically change in response to network topology changes, you should manually configure only a small number of reasonably stable routes.

**Dynamic Routes Using RIP.** Automated methods of learning routes help you keep up with a changing network environment, allowing routes to be reconfigured quickly and reliably. Interior Gateway Protocols (IGP), which operate within intranetworks, provide this automated method. The LANplex

system uses RIP (one of the most widely used IGPs), to dynamically build its routing tables.

RIP operates in terms of active and passive devices. The *active devices*, usually routers, broadcast their RIP messages to all devices in a network; they update their own routing tables when they receive a RIP message. The *passive devices*, usually hosts, listen for RIP messages and update their routing tables; they do not send RIP messages.

An active router sends a RIP message every 60 seconds. This message contains both the network number for each destination network and the number of hops to reach it. In RIP, each router that a packet must travel through to reach a destination equals one hop.

### Selecting the Best Route

Large networks have multiple routes to a single network. The routers use these criteria to select the best "route" to a network when choosing between alternate routes:

- Select the route that requires the lowest number of ticks.
- If multiple routes exist with an equal number of ticks, select the route that also has the lowest number of hops.
- If multiple routes exist with both ticks and hops equal, choose any of the routes as the "best" route.

### Service Advertising Protocol

The Service Advertising Protocol (SAP) allows servers (for example, file servers, print servers, and gateway servers) to advertise their addresses and services. Through the use of SAP, adding and removing services on an internetwork becomes dynamic. As servers are booted up, they advertise their services using SAP. When they are brought down, they use SAP to indicate that their services are no longer available.

### Internetwork Service Information

Using SAP, routers create and maintain a database of internetwork service information. Clients use this data to determine what services are available on the network and to obtain the internetwork address of the nodes (servers) where they can access desired services.



*A workstation must first know a server's network address before it can initiate a session with a file server.*

### SAP Packet Structure

SAP uses IPX and the medium-access protocols for its transport. The packet structure allows the following functions:

- A workstation request for the name and address of the nearest server of a certain type
- A router request for the names and addresses of all the servers or of all the servers of a certain type on the internetwork
- A response to a workstation or a router request
- Periodic broadcasts by servers and routers
- Changed server information broadcasts

Figure 6-5 provides an overview of the SAP packet structure. Note that the packet structure is encapsulated within the data area of IPX.

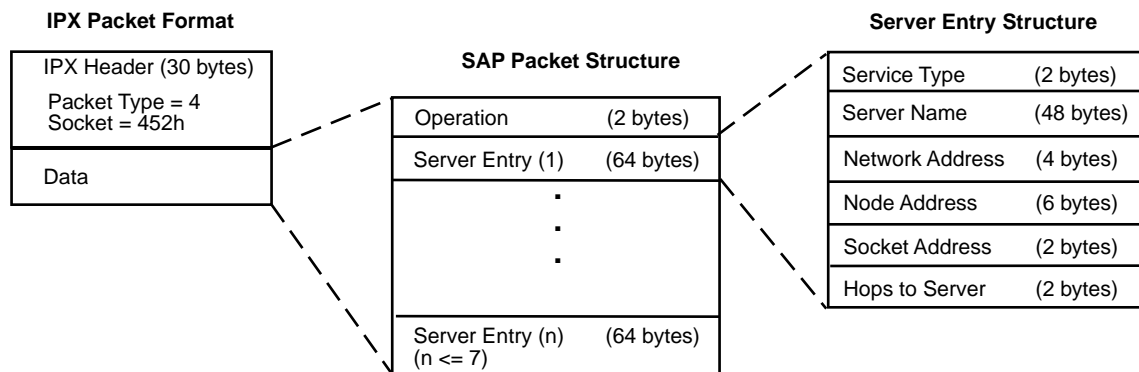


Figure 6-5 SAP Packet Structure

A SAP packet consists of the following fields:

- **Operation** — This field indicates the type of operation the SAP packet performs. It can be set to one of the following values:
  - 1=Request
  - 2=Response
  - 3=Get Nearest Server Request
  - 4=Get Nearest Server Response
- **Server Entry** — Each 64-byte server entry includes information about a particular server. It consists of the following fields:
  - **Service Type** — This 2-byte field identifies the type of service the server provides.



*Although IPX routers use SAP, routers typically do not act as servers and require no Service Type assignment.*

- **Server Name** — This field contains the 48-byte character string name that is assigned to a server. The server name, in combination with the service type, uniquely identifies a server on an internetwork.
- **Network Address** — This 4-byte field contains the server's network address.
- **Node Address** — This 6-byte field contains the server's node address.
- **Socket Address** — This 2-byte field contains the socket number that the server uses to receive service requests.
- **Hops to Server** — This 2-byte field indicates the number of intermediate networks that must be passed through to reach the server associated with this field entry. Each time the packet passes through an intermediate network, the field is incremented by 1.

By using SAP, servers can advertise their services and addresses. The information that these servers broadcast is not directly used by clients; rather it is collected by a SAP agent within each router on the server's segment. The SAP agents store this information in a server information table. If the agents reside within a server, the information is also stored in their server's bindery. The clients can then contact the nearest router or file server SAP agent for server information.



The SAP broadcasts that servers and routers send are local and, therefore, only received by SAP agents on their connected segments. However, SAP agents periodically broadcast their server information so that all SAP agents on the internetwork have information about all servers that are active on the internetwork.

### Server Information Table

A server information table holds information about all the servers on the internetwork. SAP agents use this table to store information received in SAP broadcasts. Figure 6-6 shows an example of a typical server information table.

| Server Table |         |      |          |                   |        |      |     |
|--------------|---------|------|----------|-------------------|--------|------|-----|
| Interface    | Name    | Type | Network  | Node              | Socket | Hops | Age |
| 1            | LPX1102 | 4    | 45469f33 | 00-00-00-00-00-01 | 451    | 2    | 102 |
| 1            | LPX1103 | 4    | 45469f44 | 00-00-00-00-00-01 | 451    | 5    | 65  |
| 2            | LPX2001 | 4    | 45470001 | 00-00-00-00-00-01 | 451    | 4    | 33  |

Figure 6-6 Server Information Table

The server information table provides the following information:

- **Interface** — Indicates from which interface the information was received
- **Server Name** — The name of the server
- **Server Type** — Indicates the type of service provided
- **Network Address** — The address of the network on which the server resides
- **Node Address** — The node of the server
- **Socket Address** — The socket number on which the server will receive service requests
- **Hops to Server** — The number of intermediate networks that must be passed through to reach the server associated with this entry
- **Age of Server** — The time since the last update for that server

The server information table is either statically or dynamically generated and updated.

**Static Servers.** A static server is one you manually configure in the server information table. Static servers are useful in environments where no routing protocol is used or where you want to override some of the servers generated with a routing/server protocol. Because static servers do not automatically change in response to network topology changes, you should manually configure only a small number of relatively stable servers.

**Dynamic Routes Using SAP.** An automated method of adding and removing services helps you keep up with a changing network environment, allowing servers to advertise their services and addresses quickly and reliably. SAP provides this automated method.

As servers are booted up, they advertise their services using SAP. When servers are brought down, they use SAP to indicate that their services are no longer available.

The information that these servers broadcast is not directly used by clients; rather it is collected by a SAP agent within each router on the server's segment. The SAP agents store this information in the server information table. Clients can then use the table to contact the nearest router or file server SAP agent for server information.

### Server Information Maintenance

When a router's SAP agent receives a SAP broadcast response indicating a change in the internetwork server configuration, the agent must update its server information table and inform other SAP agents of these changes. Examples of such a change are when a server is disconnected or becomes accessible through a better route.


To relay this changed information to the rest of the internetwork, the SAP agent immediately sends a broadcast to all of its directly connected segments except the segment from which the information was received. This broadcast packet contains information regarding the server change. The change information is also reflected in all future periodic broadcasts.

**SAP Aging.** Router SAP agents implement an aging mechanism to handle conditions that cause a SAP agent to go down suddenly without sending a DOWN broadcast. Examples of such changes are a hardware failure, power interruptions, and power surges. A SAP agent maintains a timer for each entry in its server information tables that keeps track of how much time has

elapsed since information was received concerning a particular table entry. Since this information is either new or changed, the SAP agent that receives this information immediately passes it on, and the change is quickly learned throughout the internetwork.

**SAP Request Handling.** When a SAP agent receives a general request, it sends the sending source a SAP response packet containing information about all servers of any type known to the receiving SAP agent. This response includes the same information sent out in a periodic broadcast. When the request is specific, the SAP agent sends a SAP response directly to the requesting node. This response contains all known information regarding all servers of the requested type.





## ROUTING IN AN APPLE TALK<sup>®</sup> ENVIRONMENT

This chapter provides an overview of AppleTalk<sup>®</sup> routing, and includes these topics:

- AppleTalk Network Elements
- AppleTalk Protocols
- About AARP

---

### About AppleTalk<sup>®</sup>

AppleTalk is a suite of protocols defined by Apple Computer, Inc., for connecting computers, peripherals devices, and other equipment on a network. AppleTalk protocols support most of the functions offered by the Open Standards Interconnect (OSI) reference model.

The AppleTalk protocols work together to provide file sharing and printer sharing, as well as applications like electronic mail and database access. All Macintosh<sup>®</sup> computers have AppleTalk connectivity options built into them, making it the de facto standard for Apple<sup>®</sup> computer networks.

---

### AppleTalk<sup>®</sup> Network Elements

An AppleTalk network consists of different nodes in groups of networks in an AppleTalk internet. These nodes can include workstations, routers, and printers, or services for other computers, called clients.

This section describes the elements of an AppleTalk internet:

- AppleTalk networks
- AppleTalk nodes
- AppleTalk zones
- Seed routers

**AppleTalk®  
Networks**

A network in an AppleTalk internet is a cable segment attached to a router. Each network is identified by a network number or range of network numbers. The network administrator assigns these numbers from a range of valid network numbers.

Two AppleTalk network numbering systems are currently in use: nonextended (Phase 1) and extended (Phase 2). 3Com routers support extended network numbers. While the LANplex system will not translate Phase 1 packets to Phase 2 packets, it will route packets to a Phase 1 network. The LANplex system anticipates that a gateway exists between the two networks to translate the packets.

An extended network can span a range of logical networks. Network numbers in an extended network consist of a range, such as 15 through 20. This numbering scheme allows for as many as 16,580,608 nodes, although the actual cables will not support this many nodes.

**AppleTalk® Nodes**

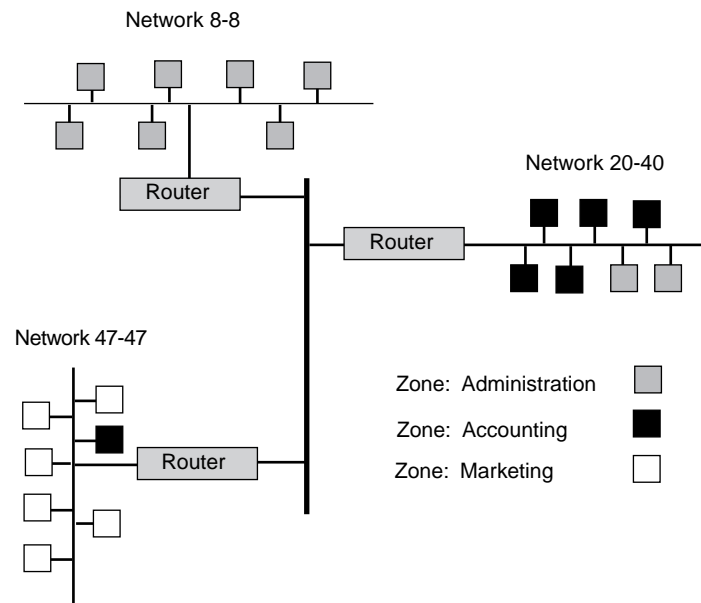
A node in a AppleTalk network is any addressable device, including workstations, printers, and routers. Nodes are physically attached to a network. Each AppleTalk node is identified by a unique AppleTalk address that each node selects at initialization. The address consists of the node's network number and a unique node number.

**Named Entities**

When a device on the network provides a service for other users, the network administrator can give the device a name. The name appears on the Chooser menu of the Macintosh with an associated icon. For example, the Chooser of the Macintosh can include a printer icon. When you select the printer icon, several printer names can appear in a list, such as Laser1, or Laser 2. The Name Binding Protocol (NBP), described later in this chapter, translates these device names into AppleTalk addresses.

**AppleTalk® Zones**

An AppleTalk zone is a logical collection of nodes on an AppleTalk internet. A zone can include all nodes in a single network or a collection of nodes in different networks. You assign a unique name to each zone to identify it in the internet. Figure 7-1 illustrates the relationship between physical AppleTalk networks and logical AppleTalk zones.



**Figure 7-1** AppleTalk Networks and AppleTalk Zones

Figure 7-1 shows an AppleTalk internet with three networks: 47-47, 20-40, and 8-8. Three AppleTalk zones span the networks in this internet: Administration, Accounting, and Marketing. Network 20-40 includes two nodes in the Administration zone and five nodes in the Accounting zone. Network 47-47 includes a node from the Accounting zone as well as the Marketing nodes. Network 8-8 consists of nodes in the Administration zone only.

Creating zones within a network reduces the amount of searching a router has to do to find a resource on the network. For example, you may want to gain access to a printer on the network. Instead of searching the whole network when you want to print a file to a certain printer, the router searches for it within a particular zone. You gain access to the printer more

quickly within the zone because the zone includes fewer devices than the entire internet does.

**Seed Routers** A seed router initializes the internet with AppleTalk configuration information, including network numbers and zone names. The seed router broadcasts this information so that nonseed routers can learn it. You can designate a seed router through the Administration Console.

A nonseed router listens for a seed router and then takes the configuration information from the first seed router it detects. After a nonseed router obtains the configuration information, it can participate in the network as if it were a seed router as well.

---

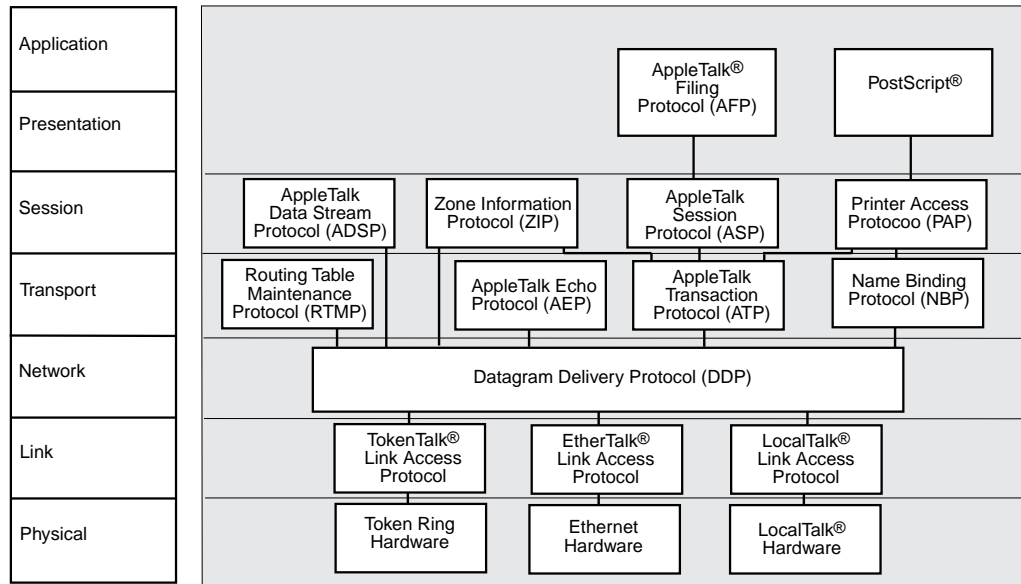
## AppleTalk Protocols

AppleTalk protocols work together to ensure the seamless flow of information throughout the AppleTalk internet. Figure 7-2 shows a simplified view of AppleTalk protocols and their relationship to the OSI reference model. Together, these protocols provide the following services:

- Physical Connectivity
- End-to-End Services
- Reliable Data Delivery



OSI Reference Model

**Figure 7-2** AppleTalk Protocols and the OSI Reference Model

The AppleTalk six-layer protocol suite is not fully compliant with the OSI seven-layer reference model. However, AppleTalk provides many of the functions and services provided by OSI. Note that AppleTalk has no specific protocols for the application layer, since the lower levels provide printer and file service.

### Physical Connectivity

The physical layer of the OSI protocol stack defines the network hardware. You can use standard network hardware, such as that defined for Ethernet and Token Ring networks, with AppleTalk. Apple has also defined its own network hardware, called LocalTalk, which uses a synchronous RS-422A bus for communications.

The data link layer provides the interface between the network hardware and the upper layers of the protocol stack. The AppleTalk data link layer includes three link access protocols (LAPs): TokenTalk LAP (TLAP), Ethernet LAP (ELAP), and LocalTalk Link Access Protocol (LLAP).

The AppleTalk Address Resolution Protocol (AARP), which translates hardware addresses to AppleTalk addresses, also exists at the datalink layer

because it is closely related to the Ethernet and token ring LAPs. This protocol is usually included in the definition of each LAP, so it does not appear in the reference model. See the section “About AARP” later in this chapter for more information about this protocol.

### **The Datagram Delivery Protocol (DDP)**

The network layer accepts data from the layers above it and divides the data into packets that can be sent over the network through the layers below it. The Datagram Delivery Protocol (DDP) transfers data in packets called datagrams.

Datagram delivery is the basis for building other AppleTalk services, such as electronic mail. The DDP allows AppleTalk to run as a process-to-process, best-effort delivery system in which the processes running in the nodes of interconnected networks can exchange packets with each other.

### **End-to-End Services**

The transport layer and the session layer provide end-to-end services in the AppleTalk network. These services ensure that routers transmit data accurately between one another. Each layer includes four protocols that work together to support these services. This section describes these protocols and provides more detail for those that you can view using the LANplex Administration Console.

#### **Transport Layer Protocols**

An AppleTalk internet has four transport layer protocols:

- Routing Table Maintenance Protocol (RTMP)
- AppleTalk Echo Protocol (AEP)
- AppleTalk Transaction Protocol (ATP)
- Name Binding Protocol (NBP)

**Routing Table Maintenance Protocol (RTMP).** The protocol maintains information about AppleTalk addresses and connections between different networks. It specifies that each router 1) learns about new routes from the other routers and 2) deletes routes after a certain period if the local router no longer broadcasts the route to the network.

Each router builds a routing table that is the basis of dynamic routing operations in an AppleTalk internet. Every 10 seconds, each router sends an RTMP data packet to the network. Routers use the information that they receive in the RTMP broadcasts to build their routing tables. Each entry in the routing table contains these items:

- The network range
- The distance in hops to the destination network
- The interface number of the destination network
- The state of each port (good, suspect, bad, really bad)

The router uses these items to determine the best path along which to forward a data packet to its destination on the network. The routing table contains an entry for each network that a datagram can reach within 15 hops of the router. The table is aged at set intervals as follows:

- 1** After a period of time, the RTMP changes the status of an entry from good to suspect.
- 2** After an additional period of time, the RTMP changes the status of an entry from suspect to bad.
- 3** After an additional period of time, the RTMP changes the status of an entry from bad to really bad.
- 4** Finally, the router removes from the table the entry of a nonresponding router with a really bad status.

The data in the routing table is cross-referenced to the Zone Information Table (ZIT). This table maps networks into zones. The section on the session layer protocols includes information about the ZIT.

Figure 7-3 illustrates a simple AppleTalk network and Table 7-1 shows the corresponding routing table.

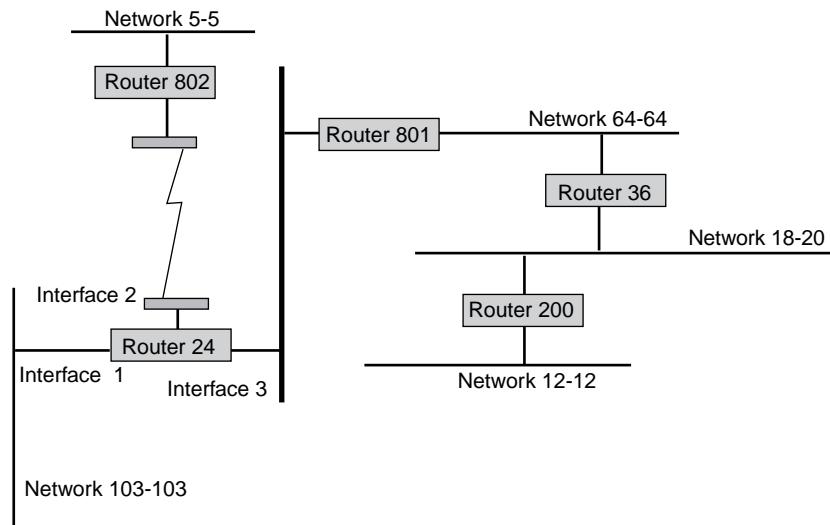


Figure 7-3 A Simple AppleTalk Network

Table 7-1 The Routing Table for Router 24 in Figure 7-3

| Network Range | Distance | Interface | State |
|---------------|----------|-----------|-------|
| 5-5           | 1        | 2         | Good  |
| 12-12         | 3        | 3         | Good  |
| 18-20         | 2        | 3         | Good  |
| 103-103       | 0        | 1         | Good  |
| 64-64         | 1        | 3         | Good  |

You can view the AppleTalk routing tables in your network through the Administration Console.

**AppleTalk Echo Protocol (AEP).** AppleTalk nodes use the AEP to send datagrams to other nodes in the network. It causes the destination node to return, or echo, the datagram to the sending node. This protocol can determine whether a node is accessible before any sessions are started, and it can enable users to estimate the round-trip delay time between two nodes.

**AppleTalk Transaction Protocol (ATP).** This protocol, along with the AppleTalk Data Stream Protocol (ADSP), ensures that DDP packets are delivered to a destination without any losses or corruption.

**Name Binding Protocol (NBP).** This protocol translates alphanumeric entity names to AppleTalk addresses. It maintains a table that references the addresses of nodes and named entities that reside in that node. Because each node maintains its own list of named entities, the names directory within an AppleTalk network is not centralized. It is a distributed database of all nodes on the internet.

### The Session Layer Protocols

An AppleTalk internet has four session-layer protocols:

- Zone Information Protocol (ZIP)
- AppleTalk Data Stream Protocol (ADSP)
- AppleTalk Session Layer Protocol (ASP)
- Printer Access Protocol (PAP)

**The Zone Information Protocol (ZIP).** ZIP works with RTMP to maintain a table that maps network numbers to network zones for the entire AppleTalk internet. Network zones are the logical groupings of AppleTalk networks. As we have seen it, the table created by ZIP is called the Zone Information Table (ZIT). The Administration Console allows you to view the zone information table by network number or network zone.

ZIP creates a zone information table in each router. Each entry in the ZIT is a "tuple," or pair, that includes a network number and a network zone name. When an NBP packet arrives at the router, it includes the zone name which the router compares with entries in the zone table. The router then matches the network number from the matching ZIT tuple to the one in the RTMP table to find the interface where it can route the packets.

**AppleTalk Data Stream Protocol (ADSP).** The ADSP works with the ATP to ensure reliable data transmission. Unlike ATP, however, ADSP provides full-duplex byte-stream delivery. This means that two nodes can communicate simultaneously. ASDP also includes flow control, so that a fast sender does not overwhelm a slow receiver.

**AppleTalk Session Protocol (ASP).** The ASP passes commands between a workstation and a server once a connection is made between the two. ASP ensures that the commands are delivered in the same order as they were sent and returns the results of these commands to the workstation.

**Printer Access Protocol (PAP).** The PAP maintains communications between a workstation and a printer or print service. The PAP functions include setting up and maintaining a connection, transferring the data, and tearing down the connection on completion of the job. Like other protocols at the session layer, PAP relies on NBP to find the addresses of named entities. PAP also depends on ATP for sending data.

**Presentation Layer** The presentation layer maintains information about files, formats, and translations between formats. An AppleTalk internet has two protocols at the presentation layer: the AppleTalk Filing Protocol (AFP) and PostScript®. AFP provides remote access to files on the network. PostScript is a paged description language used by many printers.

---

## About AARP

The AppleTalk Address Resolution Protocol (AARP) maps the hardware address of an AppleTalk node to an AppleTalk protocol address. It does this mapping for both extended and nonextended networks.

When a node on the network initializes, it randomly selects an AppleTalk address for itself. At the same time, it sends out ten AARP probe packets. The probe packets determine whether any other nodes on the network are using the address it has chosen. If a node on the network is already using that address, the node randomly selects another address and sends out another probe packet.

The AARP maintains an Address Mapping Table (AMT) with the most recently used hardware addresses and their corresponding AARP addresses. If an address is not in this table, AARP sends a request to the protocol address and adds the hardware address to the table when the destination node replies. You can view this table, called the AARP cache, through the LANplex Administration Console.





# 8


## ADMINISTERING VLANs

This chapter describes how to display information about VLANs and how to configure VLANs.

Through the Administration Console, you can:

- Display summary or detailed information on VLANs
- Define or modify a VLAN definition
- Delete a VLAN definition

### Displaying VLAN Information



You can display a summary of VLAN information or a detailed report. When you display a summary, you receive information about the protocols and ports assigned to each VLAN plus the layer 3 addresses used to manage flood domains for overlapping IP subnets. The detailed VLAN report includes the summary information plus additional utilization statistics.

From the top level of the Administration Console, enter:

**bridge vlan summary**

or

**bridge vlan detail**

The VLAN information is displayed in the format you specified.

Example of a summary display for several VLANs:

Select menu option (bridge/vlan): **summary**

| Index | Protocol | Identifier | Ports    |
|-------|----------|------------|----------|
| 1     | default  | 0          | 1-17     |
| 2     | IP       | 2          | 1, 5-7   |
| 3     | IPX      | 3          | 8-10     |
| 4     | IP       | 4          | 7, 12-15 |

| Index | Name       | Layer 3                      |
|-------|------------|------------------------------|
| 1     | none       |                              |
| 2     | eastgroup  | 158.101.111.16 255.255.255.0 |
| 3     | westgroup  | none                         |
| 4     | northgroup | 158.101.112.14 255.255.255.0 |

Example of a detailed display for the VLANs:

Select menu option (bridge/vlan): **detail**

| Index | Protocol | Identifier | Ports    |
|-------|----------|------------|----------|
| 1     | default  | 0          | 1-17     |
| 2     | IP       | 2          | 1, 5-7   |
| 3     | IPX      | 3          | 8-10     |
| 4     | IP       | 4          | 7, 12-15 |

| Index | Name       | Layer 3                      |
|-------|------------|------------------------------|
| 1     | none       |                              |
| 2     | eastgroup  | 158.101.111.16 255.255.255.0 |
| 3     | westgroup  | none                         |
| 4     | northgroup | 158.101.112.14 255.255.255.0 |

| index | inPackets | inBytes | outPackets | outBytes |
|-------|-----------|---------|------------|----------|
| 1     | 342       | 3676    | 322        | 2987     |
| 2     | 125       | 7654    | 118        | 6897     |
| 3     | 345       | 7554    | 289        | 7431     |
| 4     | 876       | 8651    | 765        | 7969     |

Table 8-1 describes these statistics.

**Table 8-1** Fields for VLAN Information

| Field      | Description   |
|------------|---|
| Index      | A system-assigned index used for identifying a particular VLAN  |
| Protocol   | The protocol suite of the VLAN  |
| Identifier | A unique, user-defined (4-byte) integer for use by global management operations                                   |
| Ports      | The numbers of the ports assigned to the VLAN   |
| Name       | A 16-byte character string intended to identify the members of the VLAN   |
| Layer 3    | Optional parameters consisting of IP subnet and mask used to set up flood domains for overlapping IP VLAN subnets |


continued

**Table 8-1** Fields for VLAN Information (continued)

| Field      | Description  |
|------------|--|
| inPackets  | Number of flooded broadcast and multicast packets that were received on the VLAN |
| inBytes    | Number of flooded broadcast and multicast bytes that were received on the VLAN   |
| outPackets | Number of flooded broadcast and multicast packets transmitted over the VLAN      |
| outBytes   | Number of flooded broadcast and multicast bytes transmitted over the VLAN        |

## Defining VLAN Information

### Top-Level Menu



Follow these steps to create a VLAN definition:

- 1 From the top level of the Administration Console, enter:  
**bridge vlan define**
- 2 Enter the appropriate protocol suite: (IP, IPX, AppleTalk, XNS, DECnet, SNA, Banyan, X.25, NetBIOS, NetBEUI, default)
- 3 Enter the VLAN interface identifier.
- 4 Enter the VLAN name, enclosed in quotation marks.
- 5 Enter the number(s) of the port(s) or **all** to assign all ports to the VLAN.

You are prompted to enter the number(s) of the port(s) that can be assigned to the VLAN.

If you did not choose the IP protocol suite for this VLAN, you have completed the steps for defining the VLAN.

If you selected the IP protocol suite, follow these steps:

- 1 Enter **defined** to use layer 3 subnet addressing and continue with steps 2 and 3, **OR** enter **undefined** to not use layer 3 addressing.
- 2 Enter the IP subnet address.
- 3 Enter the subnet mask.

Example:

```
Select menu option (bridge/vlan): define
Enter Protocol Suite
(IP,IPX,AppleTalk,XNS,DECnet,SNA,Banyan,X.25,NetBIOS,NeBEUI,
default): IP
Enter VLAN Identifier: 1
Enter VLAN Name: "SD Marketing"
Ports 1=FDDI, 2-17=Ethernet
Enter port(s) (1-17|all): 1-5
Layer 3 Address (undefined, defined): defined
Enter IP Subnet Address: 158.111.122.0
Enter subnet mask [255.255.0.0] 255.255.255.0
```

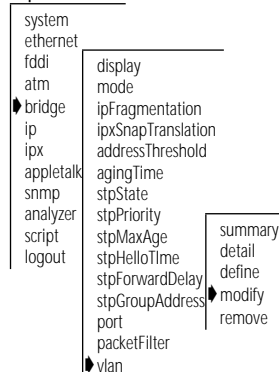


*The maximum number of VLANs you can define on a single bridge is 32.*

## Modifying VLAN Information

To modify VLAN information:

### Top-Level Menu



- 1 From the top level of the Administration Console, enter:

**bridge vlan modify**

You are prompted to reenter the information that defines the VLAN. Press the Return or Enter key to accept any value that appears in brackets [ ].

- 2 Enter the number of the VLAN interface index.

- 3 Enter the protocol suite for that VLAN: (IP, IPX, AppleTalk, XNS, DECnet, SNA, Banyan, X.25, NetBIOS, NetBEUI, default).

- 4 Enter the VLAN identifier.

- 5 Enter the VLAN name.

- 6 Enter the number(s) of the port(s) or all.

- 7 If you have selected the IP protocol suite and want to use the Layer 3 address information, enter **defined** for layer 3 addressing. Enter **undefined** if you do not want layer 3 addressing.

Example:

```
Select menu option (bridge/vlan): modify
Select VLAN interface [1-2]: 2
Protocol Suite (IP,IPX,AppleTalk,XNS,DECnet,SNA,
Banyan,X.25,NetBIOS,NetBEUI,default) [AppleTalk]: IP
VLAN Identifier [1]: 2
VLAN Name [Sales]:
Ports 1=FDDI, 2-17=Ethernet
Enter port(s) (1-17|all) [1-5]:
Layer 3 Address (undefined,defined) [undefined]:
```

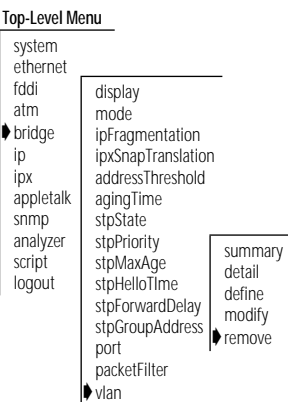
Removing VLAN Information

Follow these steps to remove a VLAN definition:

- 1 From the top level of the Administration Console, enter:  
**bridge vlan remove**
- 2 Enter the indexes for the VLANs you want to remove.

Example:

```
Select menu option (bridge/vlan): remove
Select VLAN index(es) (1-2|all): 1
```





# 9

## ADMINISTERING IP ROUTING

This chapter describes how to set up your LANplex® system to use the Internet Protocol (IP). For more information about how IP works, see Part III of this guide.

You can display or configure the following IP characteristics on your LANplex system:

- IP interfaces
- Routes
- Address Resolution Protocol (ARP) cache
- UDP Helper
- ATM ARP Server (for LANplex systems with ATM modules)
- IP Routing
- ICMP Router Discovery
- Routing Information Protocol (RIP)
- Ping
- IP statistics

---

### Administering interfaces

You can define two types of IP interfaces through LANplex Extended Switching software: IP VLAN interfaces and IP LIS interfaces. This section describes these interfaces and how to administer them.

An IP VLAN interface defines the relationship between an IP Virtual LAN (VLAN) and the subnets in the IP network. Every IP VLAN interface has one IP VLAN associated with it. Each Ethernet or FDDI switching module has one interface defined for each subnet directly connected to it. You must first define a VLAN, as described in Chapter 8, Administering VLANs, before you define an associated IP VLAN interface.

**LIS Interfaces** A logical IP subnet (LIS) interface supports logical IP over ATM. You define LIS interfaces for the ports on ATM modules only. See the Chapter 11 of the *LANplex® 2500 Operation Guide* for more information about the ATM protocol. See the *LANplex® 2500 Administration Console User Guide* for information about how to configure ATM ports.

**Interface Characteristics** Each IP interface has the following information associated with it:

- **IP Address** — This address, which is specific to your network, should be chosen from the range of addresses assigned to your organization by the central agency. This address defines both the number of the network to which the interface is attached and the interface's host number on that network.
- **Subnet Mask** — A subnet mask is a 32-bit number that uses the same format and representation as IP addresses. The subnet mask determines which bits in the IP address are interpreted as the network number, the subnet number, and the host number. Each IP address bit corresponding to a **1** in the subnet mask is in the network/subnet part of the address. Each IP address bit corresponding to a **0** is in the host part of the IP address.
- **Advertisement Address** — The switching module uses this IP address when it advertises routes to other stations on the same subnet. In particular, the system uses this address for sending RIP updates. By default the switching module uses a directed advertisement (all **1**s in the host field).
- **Cost** — This number, between 1 and 15, is used when calculating route metrics. Unless your network has special requirements, assign a cost of 1 to all interfaces.
- **Type** — The IP interface is one of these types:
  - *VLAN*, which supports routing between two VLANs
  - *LIS*, which supports classical IP over ATM
- **State** — This status of the IP interface indicates whether the interface is available for communications.
- **VLAN Interface** — When you select *VLAN* as the interface type, the Administration Console prompts you for the VLAN index number. The VLAN index number indicates which bridge ports are associated with the IP interface. When the LANplex Administration Console menu prompts you for



this option, the system displays a list of available VLAN indexes and the bridge ports associated with them.

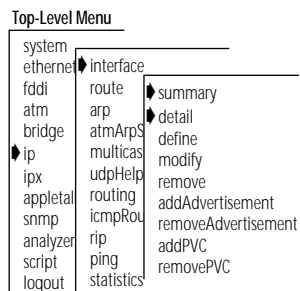
- **LIS Interface** — When you select *LIS* as the interface type, the Administration Console prompts you for LIS interface information. The information you enter depends on whether you define permanent virtual circuits (PVCs), switched virtual circuits (SVCs), or both on the LIS interface. See the *LANplex® 2500 Operation Guide* for more information on PVCs and SVCs.

If you define *SVCs*, you need to enter an ATM ARP server address. This server maintains the IP-to-ATM address translation table. You can enter the maximum number of SVCs allowed on this interface. The minimum holding time determines the least amount of time an SVC connection remains open. The inactivity timer determines how long the connection can remain open with no activity after the minimum holding time has expired. You also need to enter the ATM port number for this interface.

If you define only *PVCs* on the interface, you need to enter only the PVC numbers and the ATM port number. The other prompts do not appear because you do not enter an ATM ARP server address. If you define both SVCs and PVCs, enter all LIS interface information.

## Displaying Interfaces

You can display both summary and detailed information about all IP interfaces configured for the system. The detail display contains all the summary information as well as information about the advertisement address, PVCs, and VLANs.



To display IP interface information, enter one of the following command strings from the Administration Console top-level menu:

```
ip interface summary
```

OR

```
ip interface detail
```

## Example summary display:

IP routing is enabled, RIP is active,  
ICMP discovery is disabled.

| Index | Type | IP address  | Subnet mask   | Cost | State | VLAN | Index |
|-------|------|-------------|---------------|------|-------|------|-------|
| 1     | VLAN | 158.101.1.1 | 255.255.255.0 | 1    | Down  |      | 2     |

| Index | Type | IP address    | Subnet mask   | Cost | State | Port |
|-------|------|---------------|---------------|------|-------|------|
| 2     | LIS  | 158.101.112.1 | 255.255.255.0 | 1    | Up    | 1    |

## Example detail display:

IP forwarding is enabled, RIP is active,  
ICMP discovery is disabled.

| Index | Type | IP address  | Subnet mask   | Cost | State | VLAN | index |
|-------|------|-------------|---------------|------|-------|------|-------|
| 1     | VLAN | 158.101.1.1 | 255.255.255.0 | 1    | Down  |      | 2     |

| Index | Type | IP address    | Subnet mask   | Cost | State | Port |
|-------|------|---------------|---------------|------|-------|------|
| 2     | LIS  | 158.101.112.1 | 255.255.255.0 | 1    | Up    | 1    |

## 4 Advertisement Addresses:

158.101.112.200 158.101.112.203 158.101.112.204 158.101.112.205

## atmArpServer

47-0000-00-000000-0000-0000-00cc-080001200054-ff

| maxSvcCount | inactivityTime | minHoldingTime |
|-------------|----------------|----------------|
| 0           | 1200           | 60             |

## 1PVC:

1/32

## Defining an IP LIS Interface

When you define an IP LIS interface, you specify several general IP interface characteristics and IP LIS characteristics.

## Top-Level Menu

```

system
ethernet
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyze
  script
  logout
  interface
    summary
    detail
    define
    modify
    remove
    addAdvertisement
    removeAdvertisement
    addPVC
    removePVC
    ping
    statistics
  
```



*Before you define an IP LIS interface with SVCs, be sure you have defined an ATM ARP server as described in the section "Administering ATM ARP Servers" later in this chapter. If the LIS interface has only PVCs, you do not need to define an ATM ARP server.*

To define an IP interface:

- 1 From the top level of the Administration Console, enter:

**ip interface define**

The Console prompts you for the interface's parameters. To use the value in brackets, press [Return] at the prompt.

- 2 Enter the IP address of the interface.
- 3 Enter the subnet mask of the network to which the interface is to be connected.
- 4 Enter the cost value of the interface.
- 5 Enter the type of IP interface: LIS.
- 6 Enter the advertisement addresses for this interface. You can enter up to 32 advertisement addresses for each interface. (The maximum number on the LANplex system is 64.)
- 7 Enter the LIS information:
  - For a LIS interface with SVCs, enter the ATM ARP server address, the maximum SVC count, the inactivity timer, the minimum holding time, and the ATM port associated with the interface. (You can also accept the defaults for these values.)
  - For a LIS interface with only PVCs, enter the ATM port and the PVCs associated with the interface. You can enter up to 51 PVCs for each interface. (The maximum number on the LANplex system is 64.)

LIS interface example with both PVCs and SVCs:

```

Enter IP address: 158.101.1.1
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter cost [1]:
Enter interface type (vlan,lis) [lis]:
Enter advertisement address(es) []: 158.101.112.1
Enter ATM arp server address
[00-0000-00-000000-0000-0000-0000-000000000000-00]: 47-0000-00-000000-000
0000-00cc -000000000001-ff
Accept completed ATM address (yes,no) [yes]:
Enter max. SVC count (0=no max.0) [0]:
Enter inactivity time (0=infinite, 10-10000) seconds [1200]:
Enter min. holding time (0-10000) seconds [60]:
Select ATM port [1]:
Enter PVC(s) (VPI/VCI[]): 1/32,1/200,1/3330

```

## Defining an IP VLAN Interface

When you define an IP VLAN interface, you specify several interface characteristics, as well as the index of the VLAN associated with the interface.



*You must first define a VLAN, as described in Chapter 8, Administering VLANs, before you define an associated IP VLAN interface.*

### Top-Level Menu

|           |                     |
|-----------|---------------------|
| system    |                     |
| ethernet  |                     |
| fdi       |                     |
| atm       |                     |
| bridge    |                     |
| ip        | interface           |
| ipx       | route               |
| appletalk | arp                 |
| snmp      | atmArp              |
| analyze   | multicast           |
| script    | udpHelp             |
| logout    | routing             |
|           | icmpRo              |
|           | rip                 |
|           | ping                |
|           | statistics          |
|           | summary             |
|           | detail              |
|           | define              |
|           | modify              |
|           | remove              |
|           | addAdvertisement    |
|           | removeAdvertisement |
|           | addPVC              |
|           | removePVC           |

To define an IP VLAN interface:

- 1 From the top level of the Administration Console, enter:

```
ip interface define
```

The Console prompts you for the interface's parameters. To use the value in brackets, press [Return] at the prompt.

- 2 Enter the IP address of the interface.
- 3 Enter the subnet mask of the network to which the interface is to be connected.
- 4 Enter the cost value of the interface.
- 5 Enter the type of IP interface: VLAN.
- 6 Enter the advertisement address for this interface.
- 7 Enter the index of the VLAN associated with the interface.

Example:

```
Enter IP address: 158.101.1.1
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter cost [1]:
Enter interface type (vlan, lis) [vlan]:
Enter advertisement address(es) [158.101.1.255]:
IP VLANs:
      Index      Ports
        3        1-8
        4        9-12
Select VLAN index: 3
```



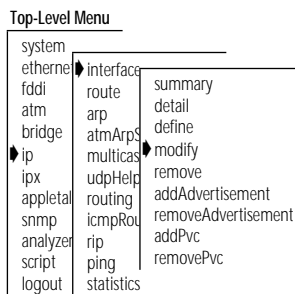
*If you physically change the configuration of your system after defining IP interfaces, the ports designated for those interfaces might no longer be valid and you might want to reconfigure your interfaces.*

## Modifying an Interface

You might want to change the configuration of an interface you have already defined.



*You can add one or more advertisement addresses or PVCs to an interface through the **addAdvertisement** and **addPVC** commands as well as through the IP interface **modify** command. If you add or change an advertisement address or PVC through the **modify** command, you must re-enter all addresses or PVCs associated with the interface, not just the one you want to add or change.*



To modify an IP interface:

- 1 From the top level of the Administration Console, enter:

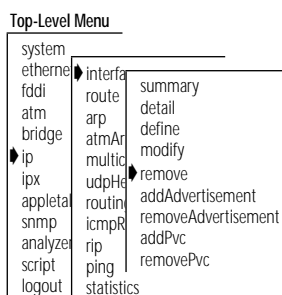
```
ip interface modify
```

You are prompted for the interface parameters. Press [Return] at the prompts for the parameters you do not want to modify.

- 2 Modify the existing interface parameters by entering a new value at the prompt.

## Removing an Interface

You might want to remove an interface if you no longer route on the ports associated with the interface.



To remove an IP interface definition:

- 1 From the top level of the Administration Console, enter:

```
ip interface remove
```

- 2 Enter the index numbers of the interfaces you want to remove.

If you have defined one or more PVCs on the interface, the Administration Console displays a message indicating that the PVCs will be removed with the interface. The following is an example of a prompt for interface 2, which has one PVC associated with it:

```
1 PVC associated with the interface index 2. Do you wish
to continue (yes/no) [yes]:
```

Accept the default (yes) if you want to delete the interface.

### Adding an Advertisement Address

This command adds an advertisement address to the advertisement address list associated with the interface.

To add an advertisement address:

- 1 From the top level of the Administration Console, enter:

```
ip interface addAdvertisement
```

- 2 Enter the interface index number.

- 3 Enter the advertisement address, separated by commas.

Example:

```
Select interface index [1]: 1
```

```
Enter advertisement address: 158.101.255.1, 158.111.1.1
```

### Removing an Advertisement Address

This command removes an advertisement address from the advertisement address list associated with the interface.

To remove an advertisement address:

- 1 From the top level of the Administration Console, enter:

```
ip interface removeAdvertisement
```

- 2 Enter the index interface number and the advertisement address you want to remove.

#### Top-Level Menu

|          |            |
|----------|------------|
| system   |            |
| ethernet | interface  |
| fdi      | route      |
| atm      | arp        |
| bridge   | atmArp     |
| ip       | multicast  |
| ipx      | udpHel     |
| applet   | routing    |
| snmp     | icmpRo     |
| analyze  | rip        |
| script   | ping       |
| logout   | statistics |


#### Top-Level Menu

|          |            |
|----------|------------|
| system   |            |
| ethernet | interface  |
| fdi      | route      |
| atm      | arp        |
| bridge   | atmArp     |
| ip       | multicast  |
| ipx      | udpHel     |
| applet   | routing    |
| snmp     | icmpRo     |
| analyze  | rip        |
| script   | ping       |
| logout   | statistics |

## Adding a Permanent Virtual Circuit (PVC)

This command adds a PVC to an LIS interface.

To add a PVC:



- 1 From the top level of the Administration Console, enter:

```
ip interface addPvc
```

- 2 Enter the index interface number that you want to associate with the PVC.
- 3 Enter the Virtual Path Interface (VPI) and the Virtual Circuit Interface (VCI) pairs in this format: **VPI/VCI**. Separate additional entries with a comma.

Example:

```
Select interface index [1]: 1
Enter [VPI/VCI]: 2/20
```

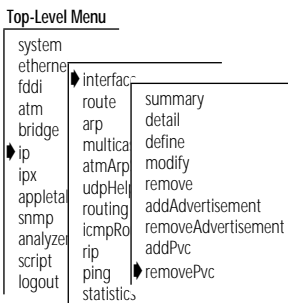
## Removing a Permanent Virtual Circuit (PVC)

This command removes one or more PVCs associated with an LIS interface.

To remove a PVC, from the top level of the Administration Console, enter:

```
ip interface removePVC
```

Enter the index number of the interface you want to remove and the VPI/VCI pair that you want to remove.



## Administering Routes

Each system maintains a table of routes to other IP networks, subnets, and hosts. You can make static entries in this table using the Administration Console or configure the system to use RIP to exchange routing information automatically.

Each routing table entry contains the following information:

- **Destination IP Address and Subnet Mask** — These elements define the address of the destination network, subnet, or host. A route matches an IP address if the bits in the IP address that correspond to the bits set in the route subnet mask match the route destination address. If the system finds

more than one routing table entry matching an address, it uses the most specific route, which is the route with the most bits set in its subnet mask. For example, the route to a subnet within a destination network is more specific than the route to the destination network.

- **Routing Metric** — This metric specifies the number of networks or subnets through which a packet must pass to reach its destination. This metric is included in RIP updates to allow routers to compare routing information received from different sources.
- **Gateway IP Address** — This address tells the router how to forward packets whose destination addresses match the route's IP address and subnet mask. The system forwards such packets to the indicated gateway.
- **Status** — For each interface, the route provides the status information in Table 9-1.

**Table 9-1** Interface Status Information

| Status     | Description                                  |
|------------|--|
| Direct     | Route goes to a directly connected network   |
| Static     | Route was statically configured              |
| Learned    | Route was learned using indicated protocol   |
| Timing out | Route was learned but is partially timed out |
| Timed out  | Route has timed out and is no longer valid   |


In addition to the routes to specific destinations, the routing table can contain an additional entry called the *default route*. The system uses the default route to forward packets that do not match any other routing table entry. You might want to use a default route in place of routes to numerous destinations that all have the same gateway IP address.



## Displaying the Routing Table

You can display a switching module's routing table to determine which routes are configured and whether the routes are operational.

### Top-Level Menu



To display the contents of the routing table, enter the following command string from the top level of the Administration Console:

**ip route display**

The example shows routes for the LANplex 2500 system. The display indicates the configuration of RIP. The default route appears as **Default Route**.

IP routing is enabled, RIP is active, ICMP router discovery is disabled.

| Destination   | Subnet mask   | Metric | Gateway     | Status        |
|---------------|---------------|--------|-------------|---------------|
| 158.101.4.0   | 255.255.255.0 | 2      | 158.101.2.8 | Static        |
| 158.101.3.0   | 255.255.255.0 | 2      | 158.101.1.2 | Learned(RIP)  |
| 158.101.2.0   | 255.255.255.  | 1      | --          | Direct        |
| 158.101.1.0   | 255.255.255.0 | 1      | --          | Direct        |
| Default Route | --            | 5      | 158.101.1.2 | Learned (RIP) |

## Defining a Static Route

Before you can define static routes, you must define at least one IP interface. Static routes remain in the table until you remove them or the corresponding interface. Static routes take precedence over dynamically learned routes to the same destination.



*Static routes are not included in periodic RIP updates sent by the system.*

To define a static route:


- 1 From the top level of the Administration Console, enter:

**ip route static**

You are prompted for the route's parameters. To use the value in brackets, press [Return] at the prompt.

- 2 Enter the destination IP address of the route.
- 3 Enter the subnet mask of the route.
- 4 Enter the gateway IP address of the route.

### Top-Level Menu

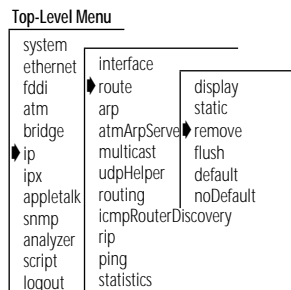


Example:

```
Enter destination IP address: 158.101.4.0
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter gateway IP address: 158.101.2.8
```

## Removing a Route

To remove a route:



- 1 From the top level of the Administration Console, enter:

```
ip route remove
```

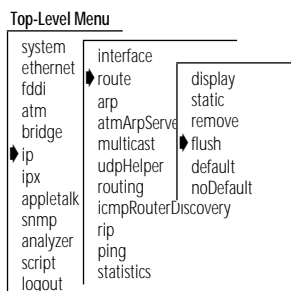
- 2 Enter the destination IP address of the route.

- 3 Enter the subnet mask of the route.

The route is immediately deleted from the routing table.

## Flushing a Route

Flushing deletes all learned routes from the routing table.



To flush all learned routes, from the top level of the Administration Console, enter:

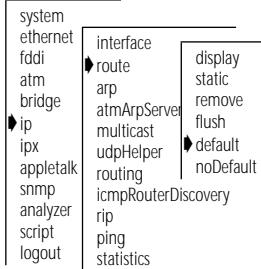
```
ip route flush
```

All learned routes are immediately deleted from the routing table.

## Setting the Default Route

If you define a default route, the system uses it to forward packets that do not match any other routing table entry. A system can learn a default route using RIP, or you can configure a default route statically.

If a system's routing table does not contain a default route — either statically configured or learned using RIP — then it cannot forward a packet that does not match any other routing table entry. If this occurs, then the module drops the packet and sends an ICMP “destination unreachable” message to the host that sent the packet.

**Top-Level Menu**

To statically configure the default route:

- 1 From the top level of the Administration Console, enter:

**ip route default**

- 2 Enter the gateway IP address of the route.

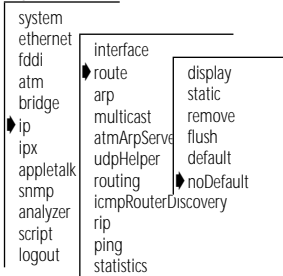
The default route is immediately added to the routing table.

### Removing the Default Route

To remove a default route, enter the following command string from the top level of the Administration Console:

**ip route noDefault**

The default route is immediately removed from the routing table.

**Top-Level Menu**

## Administering the ARP Cache

The LANplex system uses the Address Resolution Protocol (ARP) to find the MAC addresses corresponding to the IP addresses of hosts and routers on the same subnets. An ARP cache is a table of known IP addresses and their corresponding MAC addresses.

## Displaying the ARP Cache

You can display the contents of the ARP cache for your system.

To display the contents of the ARP cache, enter the following command string from the top level of the Administration Console:

```
ip arp display
```

Example display of the contents of the ARP cache:

```
IP routing is enabled, RIP is active,
ICMP router discovery is disabled
```

| IP address     | I/F | Hardware address  | Circuit |
|----------------|-----|-------------------|---------|
| 158.101.112.2  | 1   | 00-40-0b-40-64-e6 | -/-     |
| 158.101.112.7  | 1   | 08-00-20-76-a2-f2 | -/-     |
| 158.101.116.7  | 2   | 00-80-3e-02-68-00 | -/-     |
| 158.101.112.14 | 1   | 08-00-09-4e-24-20 | -/-     |
| 158.101.116.16 | 2   | 00-80-3e-02-8e-6a | -/-     |
| 158.101.116.17 | 2   | 00-80-3e-02-8e-7f | -/-     |
| 158.101.116.18 | 2   | 00-80-3e-02-8e-94 | -/-     |
| 158.101.112.22 | 1   | 08-00-20-04-d1-5e | -/-     |
| 158.101.116.19 | 2   | 00-80-3e-02-8e-2b | -/-     |
| 158.101.112.28 | 1   | 08-00-09-8c-de-3a | -/-     |
| 158.101.112.29 | 1   | 08-00-09-82-d8-1b | -/-     |
| 158.101.116.27 | 2   | 00-80-3e-1d-75-00 | -/-     |

## Removing an ARP Cache Entry

You might want to remove an entry from the ARP cache if the MAC address has changed. To remove an entry from the ARP cache:

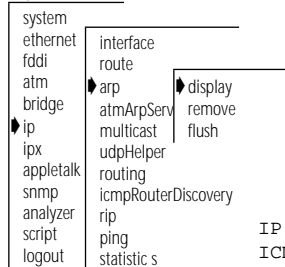
- 1 From the top level of the Administration Console, enter:

```
ip arp remove
```

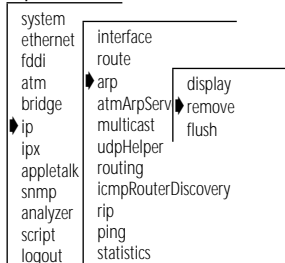
- 2 Enter the IP address you want to remove.

The address is immediately removed from the table. If necessary, the system subsequently uses ARP to find the new MAC address corresponding to that IP address.

### Top-Level Menu




### Top-Level Menu



## Flushing the ARP Cache

### Top-Level Menu



You might want to delete all entries from the ARP cache if the MAC address has changed.

To remove all entries from the ARP cache, from the top level of the Administration Console, enter:

```
ip arp flush
```

The ARP cache entries are immediately removed from the table.


## Administering ATM ARP Servers

If you are running classical IP over ATM with SVCs, you need to define an ATM ARP server for each LIS. Each LIS must connect to a single ATM network and must belong to the same IP subnet.

The **atmArpServer** menu also includes the **arp** option, which allows you to administer the ATM ARP cache.

## Displaying ATM ARP Servers

### Top-Level Menu



To display a list of ATM ARP servers, from the top level of the Administration Console, enter:

```
ip atmArpServer display
```

Example:

```
IP routing is enabled, RIP is active,
ICMP router discovery is disabled
```


| Index | Port | IP Address  | Subnet Mask   |
|-------|------|-------------|---------------|
| 1     | 1    | 158.101.1.1 | 255.255.255.0 |

```
ATM address
47-0000-00-000000-0000-0000-00cc-000000000000-ff
```

## Defining an ATM ARP Server

Determine the location of the ATM ARP server you want to use. You can define the ATM ARP server externally on another LANplex system or on an ATM switch, such as 3Com's CELLplex™ 7000 system.

### Top-Level Menu



- 1 To define an ATM ARP server, from the top level of the Administration Console, enter:  
**ip atmArpServer define**
- 2 Enter the number of the ATM port where you want to define the ATM ARP server.
- 3 Enter the IP address of the ATM port you want to define.
- 4 Enter the subnet mask. To accept the default value, shown in brackets, press the [Return] key at the prompt.

Example:


```

Select ATM port [1]
Enter IP address: 158.101.20.30
Enter subnet mask [255.255.0.0]
  
```

## Removing an ATM ARP Server

To delete a currently defined ATM ARP server, from the top level of the Administration Console, enter:

### Top-Level Menu



**ip atmArpServer remove**

The system prompts you for one or more index numbers associated with the ATM ARP servers that you want to remove. The ATM ARP server display shows the index number assigned to each ATM ARP server. The system also displays the current index numbers in the prompt.

Example:


```

Select ATM ARP server index(es) [1-2,all]: 1
  
```

## Displaying the ATM ARP Cache

To display the contents of the ATM ARP cache, from the top level of the Administration Console, enter:

### Top-Level Menu



```
ip atmArpServer arp display
```

Example:


```
IP routing is enabled, RIP is active,
ICMP router discovery is disabled
```

| IP address     | ATM Address                                    | Circuit |
|----------------|--|---------|
| 158.101.112.2  | 47-005-80-ffe100-0000-f21a-2130-80000212d0f-18 | 1/32    |
| 158.101.112.7  | 47-005-80-ffe100-0000-f22a-2130-80000211d01-18 | 1/33    |
| 158.101.116.7  | 47-005-81-ffe100-0000-f21a-2130-80000112d01-18 | 2/20    |
| 158.101.112.14 | 47-005-81-ffe100-0000-f21a-2130-80000112d01-18 | 2/22    |

## Removing an ATM ARP Cache Entry

To remove an entry from the ATM ARP cache, from the top level of the Administration Console, enter:

### Top-Level Menu



```
ip atmarpserver arp remove
```


Enter the ATM address you want to remove.

The address is immediately removed from the table.

### Flushing the ATM ARP Cache

To remove all entries from the ATM ARP cache, from the top level of the Administration Console, enter:

#### Top-Level Menu



**ip atmarpserver arp flush**

The ATM ARP cache entries are immediately removed from the table.

### Administering UDP Helper

UDP Helper allows you to send User Datagram Protocol (UDP) packets between routed networks. This protocol provides support for UDP services such as BOOTP or DHCP (Dynamic Host Configuration Protocol), that rely on the BOOTP relay agent. For example, by configuring the logical BOOTP port, you can boot hosts through the router. UDP Helper also provides a relay agent for DHCP broadcasts. UDP packets that rely on the BOOTP relay agent are modified and then forwarded through the router.

The following ports for the UDP services are mentioned in this section on UDP Helper:

- BOOTP (including DHCP) = 67
- TIME = 37
- DNS = 53

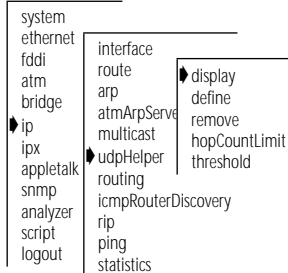
UDP Helper allows you to configure the amount of time a UDP packet is forwarded between subnetworks. UDP packets are discarded based on the hop count and the seconds value only for BOOTP and DHCP.



## Displaying UDP Helper Information

You can display the hop count and threshold configuration and list the ports with their IP forwarding addresses that are defined for your LANplex system.

### Top-Level Menu



To display UDP Helper information, enter the following command string from the top level of the Administration Console:

```
ip udpHelper display
```

Example:


```
BOOTP relay hopcount limit is 4, BOOTP relay threshold is 0.
```

```
UDP port      forwarding address
67            <158.101.1.112
```

## Defining a Port and an IP Forwarding Address

You can define port numbers and IP forwarding addresses for the UDP Helper. You may have up to 32 combinations of port numbers/IP forwarding addresses per router. You may also have multiple IP address entries for the same ports.

### Top-Level Menu



To define port numbers and IP forwarding addresses:

- 1 From the top level of the Administration Console, enter:
- 2 Enter the port numbers and IP forwarding addresses you want to define.

```
ip udpHelper define
```

## Removing a Port or an IP Forwarding Address

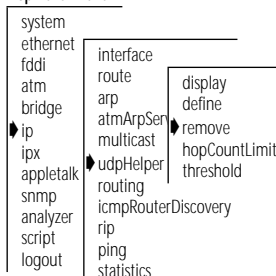
You can remove a port number or IP forwarding address defined for UDP Helper.

To remove a port number or IP forwarding address:

- 1 From the top level of the Administration Console, enter:
- 2 Enter the UDP port number that you want to remove.
- 3 Enter the IP forwarding address that you want to remove.

The port numbers and IP forwarding addresses you specified are immediately removed.


### Top-Level Menu



## Setting the BOOTP Hop Count Limit

You can set the maximum hop count for a packet to be forwarded through the router. The range is 0 through 16. The default is 4.

### Top-Level Menu



To set the hop count limit:

- 1 From the top level of the Administration Console, enter:

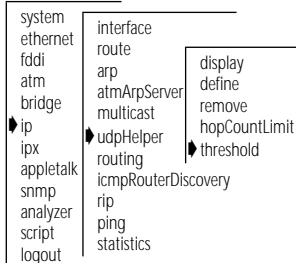
```
ip udpHelper hopCountLimit
```

- 2 Enter the BOOTP relay hop count limit.

## Setting the BOOTP Relay Threshold

You can set the maximum time limit that a packet is forwarded through the router. If you use 0 as threshold value, the router ignores the seconds field. If you use a non zero value, the router uses that value along with the hop count value to determine whether to forward the UDP packet.

### Top-Level Menu



To set the BOOTP relay threshold:

- 1 From the top level of the Administration Console, enter:

```
ip udpHelper threshold
```

- 2 Enter the BOOTP relay threshold value.

## Enabling and Disabling IP Routing

You can control whether the system forwards or discards IP packets addressed to other hosts. When you enable IP routing, the switching module acts as a normal IP router, forwarding IP packets from one subnet to another when required. When you disable IP routing, the system discards any IP packets not addressed directly to one of its defined IP interfaces.

By default, IP routing is *disabled* on the system.


To enable or disable IP routing:

- 1 From the top level of the Administration Console, enter:

```
ip routing
```

- 2 Enter the IP routing state (**enable** or **disable**).

### Top-Level Menu



## Enabling and Disabling ICMP Router Discovery

The Internet Control Message Protocol (ICMP) Router Discovery protocol (RFC 1256) allows an appropriately configured end station to locate one or more routers on the LAN to which it is attached. The end station then automatically installs a default route to each of the routers running ICMP Router Discovery. You do not need to manually configure a default route. While IP traffic may initially be directed to any of the routers on the LAN, ICMP redirect messages will subsequently channel the IP traffic to the correct router.

Only certain end stations, such as Solaris® workstations, can be configured to work with the ICMP Router Discovery protocol. Refer to the documentation for your workstation to determine whether you can configure it to work with this protocol.

To enable ICMP Router Discovery, from the top level of the Administration Console, enter

```
ip icmpRouterDiscovery
```

Enter the ICMP Router Discovery mode (**enabled** or **disabled**). This protocol is *disabled* by default.

### Top-Level Menu

|           |                     |
|-----------|---------------------|
| system    | interface           |
| ethernet  | route               |
| fdi       | arp                 |
| atm       | atmArpserver        |
| bridge    | multicast           |
| ip        | udpHelper           |
| ipx       | routing             |
| appletalk | icmpRouterDiscovery |
| snmp      | rip                 |
| analyzer  | ping                |
| script    | statistics          |
| logout    |                     |

## Setting the RIP Mode

You can select a RIP mode that is appropriate for your network. RIP can operate in any of three modes:

- **Off** — The station ignores all incoming RIP packets and does not generate any RIP packets of its own.
- **Active** — The station processes all incoming RIP packets, responds to explicit requests for routing information, and broadcasts periodic and triggered RIP updates.
- **Passive** — The station processes all incoming RIP packets and responds to explicit requests for routing information, but it does *not* broadcast periodic or triggered RIP updates.

*RIP default mode* By default, RIP operates in *passive* mode.

To set the RIP operating mode:

| Top-Level Menu |                     |
|----------------|---------------------|
| system         | interface           |
| ethernet       | route               |
| fddi           | arp                 |
| atm            | atmArpServer        |
| bridge         | multicast           |
| ip             | udpHelper           |
| ipx            | routing             |
| appletalk      | icmpRouterDiscovery |
| snmp           | rip                 |
| analyzer       | ping                |
| script         | statistics          |
| logout         |                     |

- 1 From the top level of the Administration Console, enter:

```
ip rip
```

- 2 Enter the RIP mode (**off**, **passive**, or **active**). To use the value in brackets, press [Return] at the prompt.

Example:

```
Select RIP mode (off,passive,active) [passive]: active
```

## Pinging an IP Station

Pinging uses the Internet Control Message Protocol (ICMP) echo facility to send an ICMP echo request packet to the IP station you specify. It then waits for an ICMP echo reply packet. Possible responses from pinging:

- Alive
- No answer
- Network is unreachable

A network is unreachable when there is no route to that network.

To ping an IP station:

| Top-Level Menu |                     |
|----------------|---------------------|
| system         | interface           |
| ethernet       | route               |
| fddi           | arp                 |
| atm            | atmArpServer        |
| bridge         | multicast           |
| ip             | udpHelper           |
| ipx            | routing             |
| appletalk      | icmpRouterDiscovery |
| snmp           | rip                 |
| analyzer       | ping                |
| script         | statistics          |
| logout         |                     |

- 1 From the top level of the Administration Console, enter:

```
ip ping
```

- 2 Enter the IP address of the station you want to ping.

```
IP Address: 192.9.200.40
```

You may receive one of the following responses:

```
192.9.200.40 is alive
```

OR

```
no answer from 192.9.200.40
```

For a remote IP address, you can also receive the following response:

```
Network is unreachable
```

## Displaying IP Statistics

### Top-Level Menu

|           |                     |
|-----------|---------------------|
| system    | interface           |
| ethernet  | route               |
| fdi       | arp                 |
| atm       | atmArpServer        |
| bridge    | multicast           |
| ip        | udpHelper           |
| ipx       | routing             |
| appletalk | icmpRouterDiscovery |
| snmp      | rip                 |
| analyzer  | ping                |
| script    | statistics          |
| logout    |                     |

To display IP statistics, enter the following from the top level of the Administration Console:

**ip statistics**

Example:

IP routing is enabled, RIP is active, ICMP router discovery is disabled.

|            |               |             |              |
|------------|---------------|-------------|--------------|
| inReceives | forwDatagrams | inDelivers  | outRequests  |
| 51213      | 49743         | 3227        | 2285         |
|            | outNoRoutes   | inHdrErrors | inAddrErrors |
|            | 273           | 7           | 0            |

Table 9-2 describes the IP statistics.

**Table 9-2** IP Statistics

| Field         | Description   |
|---------------|---|
| inReceives    | Total number of IP datagrams received, including those with errors  |
| forwDatagrams | Number of datagrams that the IP station attempted to forward  |
| inDelivers    | Number of datagrams that the IP station delivered to local IP client protocols                                |
| outRequests   | Number of datagrams that local IP client protocols passed to IP for transmission.                             |
| outNoRoutes   | Number of datagrams that the IP station discarded because there was no route to the destination               |
| inHdrErrors   | Number of datagrams that the IP station discarded because the IP header contained errors                      |
| inAddrErrors  | Number of datagrams that the IP station discarded because of an error in the source or destination IP address |



# 10

## ADMINISTERING IP MULTICAST ROUTING

This chapter describes how to set up your LANplex® system to use IP multicast routing. You should have previously defined IP interfaces and routes as described in Chapter 9: Administering IP Routing, before you define any IP multicast interfaces.

This appendix includes information on how to display or configure the following parameters:

- Enabling and disabling the Distance Vector Multicast Routing Protocol (DVMRP)
- Enabling and disabling the Internet Group Membership Protocol (IGMP)
- Administering IP multicast interfaces
- Administering multicast tunnels
- The Route display
- The Cache display



*To use IP multicast routing on the LANplex system, you must have already defined one or more IP interfaces. See Chapter 9, Administering IP Routing.*

## Enabling and Disabling DVMRP

DVMRP is the simple Distance Vector Multicast Routing Protocol, similar to the IP Routing Information Protocol. Multicast routers exchange distance vector updates that contain lists of destinations and the distance in hops to each destination. The routers maintain this information in a routing table.

To run multicast routing, you must enable DVMRP, which enables DVMRP on all IP interfaces that have not been disabled.

To enable or disable DVMRP, from the top level of the Administration Console, enter:

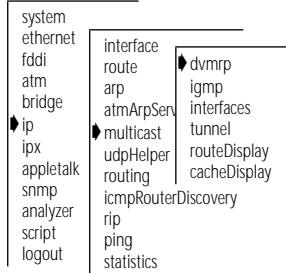
```
ip multicast dvmrp
```

The interface prompts you to enable or disable DVMRP. The default is *disabled*.

Example:

```
dvmrp mode (enabled/disabled)[disabled]: enabled
```

### Top-Level Menu

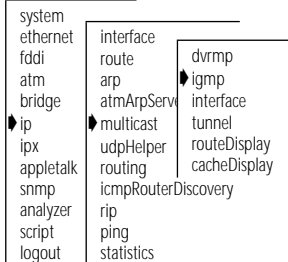


## Enabling and Disabling IGMP

IGMP enables a router or switch to determine whether group members exist in a subnetwork, or "leaf," of the Spanning Tree. It uses two search methods to make this determination:

- **Query mode** — The router or switch with the lowest IP address in the LAN broadcasts a query to all other members of the subnetwork to determine whether they are also in the group. End-stations respond to the query with IGMP packets, which report the multicast group to which they belong.
- **Snooping mode** — A router or switch performs dynamic multicast filtering based on IGMP snooping. This procedure ensures that multicast packets are flooded only to the appropriate ports within a routing interface.



**Top-Level Menu**

When you select the IGMP option, the interface prompts you to enable or disable IGMP snooping mode and IGMP query mode. Both are *enabled* by default. Under most conditions, IGMP snooping mode and IGMP query mode should remain enabled.

To enable or disable IGMP, from the top level of the Administration Console, enter:

```
ip multicast igmp
```

The interface prompts you to enable or disable IGMP query mode and IGMP snooping mode.

Example:

```
Enter igmp snooping mode
(enabled/disabled) [enabled]: enabled
Enter igmp query mode (enabled/disabled) [enabled]: enabled
```

## Administering IP Multicast Interfaces

The IP multicast interface selections allow you to enable and disable multicast characteristics on previously defined IP interfaces. A multicast interface has three characteristics, explained next.

### DVMRP Metric Value

The DVMRP metric value determines the cost of a multicast interface. The higher the cost, the less likely it is that the packets will be routed over the interface. The default value is *1*.

### Time To Live (TTL) Threshold

The TTL threshold determines whether the interface will forward multicast packets to other switches and routers in the subnetWORK. If the interface TTL is greater than the packet TTL, then the interface does not forward the packet. The default value is *1*, which means that the interface will forward all packets.


## Rate Limit

The rate limit determines how fast multicast traffic can travel over the interface in kilobytes per second. Multicast traffic may not exceed this rate limit or the LANplex system will drop packets in order to maintain the set rate. The default is set to 0, which implies no rate limit. In all other instances, the lower the rate limit, the more limited the traffic over the interface.

## Displaying Multicast Interfaces

To display a multicast interface:

### Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
ip multicast interface display
```

- 2 Enter the index numbers of the interfaces you want to display.

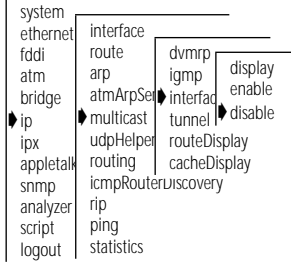
Example multicast interface configuration:

| Index | Local Address  | Metric | Threshold | RateLimit       | State        |
|-------|----------------|--------|-----------|-----------------|--------------|
| 1     | 158.101.112.32 | 1      | 1         | 0               | queries      |
|       |                | pkts   | in:0      | pkts            | out:0        |
|       | port           | 3      | peers     | 158.101.112.204 | (3.6) (0x8e) |
|       |                |        |           | 158.101.112.202 | (3.6) (0x8f) |
|       | port           | 3      | groups    | 224.2.127.255   | (3.6) (0x8e) |
|       |                |        |           | 224.2.143.24    |              |
|       | port           | 4      | groups    | 224.2.143.24    |              |
|       |                |        |           | 224.2.127.225   |              |

## Disabling Multicast Interfaces

To disable multicast routing on an interface:

### Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
ip multicast interface disable
```

- 2 Enter the index number of the interface you want to disable.


The interface is disabled.

## Enabling Multicast Interfaces

Multicast routing is enabled on all existing IP interfaces when you have not specifically disabled it. You can use this option to change the characteristics of an existing interface or to enable an interface that you had previously disabled.

To enable a multicast interface or modify the multicast characteristics of an existing IP interface:

### Top-Level Menu



- 1 From the top level of the Administration console, enter:

```
ip multicast interface enable
```

- 2 Enter the index number(s) of the interface(s) you want to enable.

- 3 Enter the DVMRP metric value of the chosen interface(s).

- 4 Enter the Time To Live (TTL) threshold of the chosen interface(s).

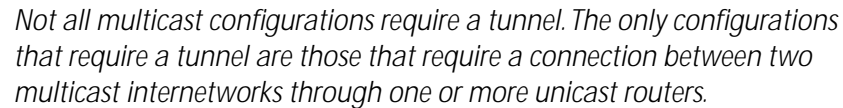
- 5 Enter the rate limit of the chosen interface(s).

Example:

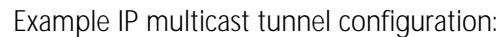
```

Enter an IP interface index [1]: 2
Enter Interface DVMRP metric [1]: 1
Enter Interface TTL threshold [1]:
Enter interface rate limit in KBits/sec [0]:
  
```

A multicast tunnel allows multicast packets to cross several unicast routers to a destination router that supports multicast. A tunnel has two end points. The local end point is associated with an interface on the LANplex router.



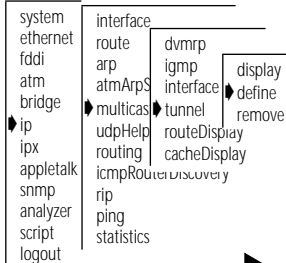
To display the IP multicast tunnel(s) on the router, from the top level menu of the Administration Console, enter:



| Index | Local Address   | Remote Address | Metric        | Threshold | RateLimit | State |
|-------|-----------------|----------------|---------------|-----------|-----------|-------|
| 1     | 158.101.112.204 | 137.39.229.98  | 2             | 255       | 500       |       |
|       |                 | pkts in:320069 | pkts out:0    |           |           |       |
|       |                 | peers          | 137.39.229.98 | (3.8)     | (0xe)     |       |

## Defining a Multicast Tunnel

### Top-Level Menu



To define an IP multicast tunnel:

- 1 From the top level of the Administration Console, enter:  
**ip multicast tunnel define**
- 2 Enter the index number(s) of the interface(s) with which you want to associate a multicast tunnel.
- 3 Enter the IP address of the destination multicast router.



*The IP address of the destination multicast router must be a remote address. The destination router cannot be directly connected to the same subnetworks as the local IP address.*

- 4 Enter the DVMRP metric value of the tunnel.
- 5 Enter the Time To Live (TTL) threshold of the tunnel.
- 6 Enter the rate limit of the tunnel.


Example:

```

Enter an IP interface index [1]: 2
Enter remote IP address: 192.9.200.40
Enter tunnel DVMRP metric [1]: 1
Enter tunnel TTL threshold [1]:
Enter tunnel rate limit [0]:
  
```

## Removing a Multicast Tunnel

### Top-Level Menu



To remove an IP multicast tunnel:

- 1 From the top level of the Administration Console, enter:  
**ip multicast tunnel remove**
- 2 Enter the index number(s) of the interfaces associated with the tunnel you want to remove.

Example:

```

Enter an IP interface index [1]: 2
  
```

The tunnel is removed.

## Displaying Routes

To display all available routes in the IP multicast routing table:

### Top-Level Menu

|           |                     |
|-----------|---------------------|
| system    |                     |
| ethernet  | interface           |
| fdi       | route               |
| atm       | arp                 |
| bridge    | atmArpSe            |
| ip        | multicast           |
| ipx       | udpHelp             |
| appletalk | routing             |
| snmp      | icmpRouterDiscovery |
| analyzer  | rip                 |
| script    | ping                |
| logout    | statistics          |

- 1 From top level of the Administration Console, enter:

**ip multicast routeDisplay**

The DVMRP status and IGMP status appear on the screen.

The following display shows all available multicast routes:

| Multicast Routing Table (2598 entries) |               |        |     |       |         |  |
|--|---------------|--------|-----|-------|---------|--|
| Origin-Subnet                          | From-Gateway  | Metric | Tmr | In-If | Out-Ifs |  |
| 157.88.29.1/32                         | 137.39.229.98 | 18     | 25  | T1    | I1      |  |
| 137.39.2.254/32                        | 137.39.229.98 | 5      | 25  | T1    | I1      |  |
| 131.215.125.236/32                     | 137.39.229.98 | 14     | 25  | T1    | I1      |  |
| 130.118.106.254/32                     | 137.39.229.98 | 10     | 25  | T1    | I1      |  |
| 129.127.118.12/32                      | 137.39.229.98 | 10     | 25  | T1    | I1      |  |
| 129.127.110.12/32                      | 137.39.229.98 | 10     | 25  | T1    | I1      |  |
| 129.127.110.11/32                      | 137.39.229.98 | 13     | 25  | T1    | I1      |  |
| 129.127.110.5/32                       | 137.39.229.98 | 10     | 25  | T1    | I1      |  |
| 129.95.63.12/32                        | 137.39.229.98 | 13     | 25  | T1    | I1      |  |
| 129.95.63.11/32                        | 137.39.229.98 | 31     | 25  | T1    | I1*     |  |
| 129.95.63.9/32                         | 137.39.229.98 | 13     | 25  | T1    | I1      |  |
| 129.95.63.8/32                         | 137.39.229.98 | 13     | 25  | T1    | I1      |  |
| 129.95.63.6/32                         | 137.39.229.98 | 13     | 25  | T1    | I1      |  |
| 129.95.63.2/32                         | 137.39.229.98 | 13     | 25  | T1    | I1      |  |
| 129.95.48.4/32                         | 137.39.229.98 | 13     | 25  | T1    | I1      |  |
| 129.95.48.3/32                         | 137.39.229.98 | 13     | 25  | T1    | I1      |  |
| 129.95.48.2/32                         | 137.39.229.98 | 13     | 25  | T1    | I1      |  |

Table 10-1 describes the fields in the route display.

**Table 10-1** Field Attributes for Multicast route display

| Field                         | Description   |
|-------------------------------|---|
| Origin-Subnet                 | The source address and the number of bits in the subnetwork   |
| From-Gateway                  | The interface address of the gateway  |
| Metric                        | The hop count   |
| Tmr                           | The amount of time, in seconds, since the routing table entry was last reset  |
| In-If <sup>1</sup>            | Interface number on which that gateway is connected. Traffic is expected to originate on this interface.<br><br>T represents the tunnel; P denotes that a prune has been sent to this tunnel. |
| Out-If <sup>1</sup>           | Set of interfaces on which the traffic will be flooded out. I represents the interface.   |
| <sup>1</sup> In-If and Out-If | Together, these attributes define a Spanning Tree configuration. Interfaces that comprise loops are disabled  |

## Displaying the Multicast Cache

The IP multicast cache contains the IP source address and destination address for packets observed on the system. The multicast cache shows you how information is routed over interfaces and ports in your system.

To display all learned routes in the multicast cache:

- 1 From the top level of the Administration Console, enter:

**ip multicast cacheDisplay**

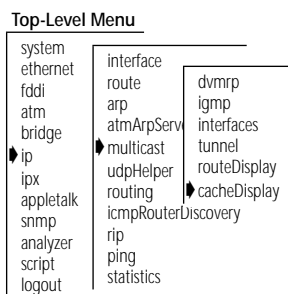
You are prompted for the multicast source address.

- 2 Enter the multicast source subnetwork address.

You are prompted for the multicast group address.

- 3 Enter the multicast group address.

The DVMRP status and IGMP status appear on the screen.



Example:

```
Enter multicast source address [131.188.0.0]
Enter multicast group address [244.2.0.2]
```

DVMRP is enabled, IGMP snooping is enabled

The following display shows the multicast cache configuration:

```
Multicast Routing Cache Table (125 entries)
  Origin                Mcast-group      CTmr  Age  PTmr  In-If      Out-Ifs
>202.242.133.128/26 224.2.0.1      7m   11m   6m  T1P        I1p
  202.242.133.139    2 packets
>128.84.247/24      224.2.0.1      2m   36m   2m  T1P        I1p
  128.84.247.53      43 packets
  128.84.247.156     33 packets
>128.138.213/24      224.2.0.1      3m    2h   2m  T1P        I1p
  128.138.213.1      23 packets
>128.206.212/24      224.2.0.1     92s   36m  60s  T1P        I1p
  128.206.212.69     8 packets
>131.136.234/24      224.2.0.1      3m   57m   3m  T1P        I1p
  131.136.234.103    12 packets
>138.39.25/24        224.2.0.1    103s    4h   71s  T1P        I1p
  138.39.25.48       46 packets
>192.5.28/24         224.2.0.1     80s    2h  48s  T1P        I1p
  192.5.28.43       178 packets
>199.94.220/24        224.2.0.1    104s    1h   72s  T1P        I1p
  199.94.220.184     10 packets
>199.104.80/24        224.2.0.1      3m   32m   3m  T1P        I1p
  199.104.80.5        4 packets
>132.197.248/21      224.2.0.1      4m    6m   4m  T1P        I1p
  132.197.248.20     1 packets
>131.188/16          224.2.0.1      3m    5h   3m  T1P        I1p
  131.188.2.54       *2492 packets 184408 bytes
>149.127/16          224.2.0.1      2m    5h  90s  T1P        I1p
  149.127.6.181      56 packets
```




Table 10-2 describes the fields in the cache display.

**Table 10-2** Information in the cache display

| Field       | Description   |
|-------------|---|
| Origin      | The source of the incoming packets. Entries preceded by an angle bracket (>) indicate a multicast subnetwork. Entries without an angle bracket, beneath the subnetwork entries, are multicast routers within that subnetwork. |
| Mcast-group | The destination multicast group   |
| CTmr        | Cache timer. The amount of time in seconds (s), minutes (m), and hours (h), that a cache entry has to remain in the cache   |
| Age         | Number of seconds (s), minutes (m), or hours (h) that the cache entry has existed.  |
| PTmr        | The time remaining, in seconds (s), minutes (m), or hours (h), before another prune will be sent to prune the Spanning Tree.  |
| In-If       | Interface number on which that gateway is connected. Traffic is expected to originate from this interface.<br><br>T represents the tunnel; P denotes that a prune has been sent to this tunnel.                               |
| Out-If      | Set of interfaces on which the traffic will be flooded out. I represents the interface.   |



## CHAPTER 10: ADMINISTERING IP MULTICAST ROUTING



# ADMINISTERING IPX ROUTING

This chapter describes how to set up your LANplex® system to use the Internet Packet Exchange (IPX) protocol to route packets. For more information about how IPX works, see Part III of this Guide.

You can display and configure the following on your LANplex system:

- IPX interfaces
- Routes
- Servers
- IPX forwarding
- Routing Information Protocol (RIP)
- Enhanced RIP mode
- Service Advertising Protocol (SAP)
- IPX statistics

---

## Administering Interfaces

An IPX interface defines the relationship between an IPX Virtual LAN (VLAN) and the IPX network. Every IPX interface has one IPX VLAN associated with it. Each switching module has one IPX interface defined for each subnet directly connected to it. You must first define a VLAN, as described in Chapter 8: Administering VLANs, before you define an associated interface.

An IPX interface has the following information associated with it:

- **IPX network address** — The network administrator sets this 4-byte address. Each address within the network should be unique.
- **Cost** — This number, between 1 and 15, is used when calculating route metrics. Unless your network has special requirements, such as the need for redundant paths, you should assign a cost of **1** to each interface.
- **Encapsulation format** — IPX routing uses four Ethernet encapsulation formats and two FDDI encapsulation formats. The Ethernet encapsulation formats are Ethernet Type II, Novell 802.3 raw, 802.2 LLC, and 802.3 SNAP. The FDDI encapsulation formats are FDDI 802.2 and FDDI SNAP.

The two FDDI encapsulation formats correspond to the Ethernet 802.2 LLC and 802.3 SNAP encapsulation formats. If you select either of these Ethernet encapsulation formats, the corresponding FDDI encapsulation format is automatically selected for shared Ethernet and FDDI ports.

- **State** — The status of the IPX interface indicates whether the interface is available for communications (*Up*) or unavailable (*Down*).
- **VLAN index** — The VLAN index indicates which bridge ports are associated with the IPX interface. When the interface prompts you for this option, it displays a list of available VLAN indexes and the ports associated with them.

## Displaying IPX Interfaces

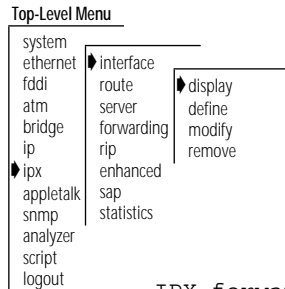
You can display a table that shows all IPX interfaces and their parameter settings configured for the system.

To display IPX interface information:

From the Administration Console top-level menu, enter:

**ipx interface display**

As shown in the following example, the current configuration is displayed. It contains IPX forwarding, RIP, and SAP information for the system as well as IPX interface information.



IPX forwarding is enabled, RIP is active, SAP is active.

| Index | IPX address | Cost | Format | State | VLAN index |
|-------|-------------|------|--------|-------|------------|
| 1     | 45469f30    | 1    | 802.2  | Up    | 2          |
| 2     | 5d41a110    | 1    | 802.2  | Up    | 1          |
| 3     | 6d321a22    | 1    | 802.2  | Up    | 4          |

## Defining an IPX Interface

When you define an interface, you define the interface's IPX address, cost, format, and the associated IPX VLAN index.



*You must define an IPX VLAN before you define the IPX interface to associate with that VLAN.*

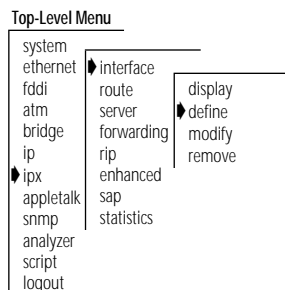
To define an IPX interface:

- 1 From the Administration Console top-level menu, enter:

**ipx interface define**

You are prompted for the interface's parameters. To use the value in brackets, press [Return] at the prompt.

- 2 Enter the IPX network address of the interface.
- 3 Enter the cost of the interface.
- 4 Enter the format of the interface.
- 5 Enter the index of the IPX VLAN associated with this interface.



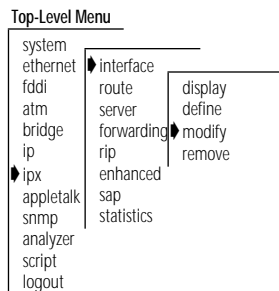
Example:

```
Enter IPX Address: 0x45469f30
Enter Cost [1]: 1
Enter Frame Format (Ethernet II: 0, 802.2: 1, Raw 802.3: 2, SNAP: 3): 1
IPX VLANs:
      Index      Ports
        3         1-8
        4         9-12
Select VLAN index: 3
```

### Modifying an Interface

You might want to change the configuration of an interface that you have already defined.

To modify an IPX interface:



- 1 From the Administration Console top-level menu, enter:

**ipx interface modify**

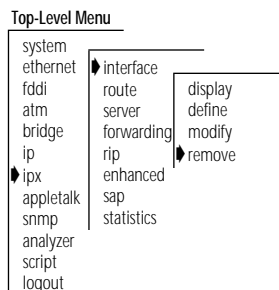
You are prompted for the interface parameters. Press [Return] at the prompts for which you do not want to modify the value.

- 2 Modify the existing interface parameters by entering a new value at the prompt.

### Removing an Interface

You may want to remove an interface if you no longer perform routing on the ports associated with the interface.

To remove an IPX interface definition:



- 1 From the Administration Console top-level menu, enter:

**ipx interface remove**

- 2 Enter the index number(s) of the interface(s) you want to remove.

The interface is removed.

---

## Administering Routes

Your system maintains a table of routes to other IPX networks. You can either use the Routing Information Protocol (RIP) to exchange routing information automatically or make static entries in this table using the Administration Console.

Each routing table entry contains the following information:


- **Address** — The 4-byte IPX network address of a segment currently known to the router.
- **Hops** — The number of routers that must be crossed to reach the network segment. The maximum number of routers a packet can cross is 15. Exception: An IPX NetBIOS packet can cross no more than 7 routers.
- **Tics** — An estimate of how long it will take the packet to reach this segment. A tic is approximately 55 milliseconds.
- **Node** — The 6-byte address of the router that can forward packets to the segment. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.
- **Age** — The number of seconds that have elapsed since the last time the route was heard from.

## Displaying the Routing Table

You can display the routing tables for the system to determine which routes are configured and if they are operational.

To display the contents of the routing table, from the Administration Console top-level menu, enter:

### Top-Level Menu



**ipx route display**

The example displays the configuration of IPX forwarding, RIP, and SAP, as well as the routing table.

IPX forwarding is enabled, RIP is active, SAP is active


| Interface | Address  | Hops | Tics | Node              | Age |
|-----------|----------|------|------|-------------------|-----|
| 2         | 45469f02 | 5    | 6    | 08-00-02-04-80-b6 | 44  |
| 2         | c2c028ca | 4    | 28   | 08-00-02-04-80-b6 | 85  |
| 2         | aaaaaaaa | 6    | 671  | 08-00-02-04-80-b6 | 85  |

## Defining a Static Route

Before you define static routes on the system, you must define at least one IPX interface. See the section “Defining an IPX Interface” on page 11-3. Static routes remain in the table until you remove them or until you remove the corresponding interface. Static routes take precedence over dynamically learned routes to the same destination. You can set up to 16 static routes.

To define a static route:

### Top-Level Menu



- 1 From the Administration Console top-level menu, enter:

**ipx route static**

- 2 Enter the 4-byte IPX network address of the route.
- 3 Enter the cost of the route.
- 4 Enter the interface number of the route.




- 5 Enter the node address of the route.

A static route is defined in the following example:

```
Enter IPX address: 0x45469f30
Enter Cost: 1
Enter Interface number: 1
Enter node address: 08-00-3e-22-15-78
```

### Removing a Route

To remove a route:



- 1 From the Administration Console top-level menu, enter:

```
ipx route remove
```


- 2 Enter the 4-byte IPX network address.

The route is immediately deleted from the routing table.

### Flushing Routes

Flushing deletes all dynamically learned routes from the routing table.

To flush all learned routes from the Administration Console top-level menu, enter:



```
ipx route flush
```

All learned routes are immediately deleted from the routing table.

---

## Administering Servers

Your system maintains a table of servers that reside on other IPX networks. You can either use the Service Advertising Protocol (SAP) to exchange server information automatically or make static entries in this server table using the Administration Console.

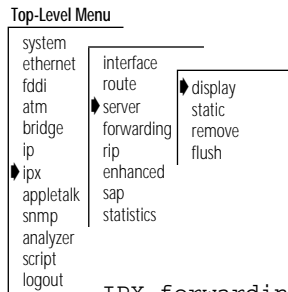
Each server table contains the following information:

- **Name** — The user-defined name of the server.
- **Type** — The type of service provided by the server.
- **Node** — The 6-byte address of the server that can forward packets to the segment.
- **Socket** — The 2-byte socket address on which the server will receive service requests.
- **Hop** — The number of networks that must be crossed to reach the server. The maximum number is fifteen.
- **Age** — The number of seconds that have elapsed since the last time a server in the table was heard from.

## Displaying the Server Table

You can display the server table for the system to determine which servers are learned and if they are operational.

To display the contents of the server table, from the Administration Console top-level menu, enter:



**ipx server display**

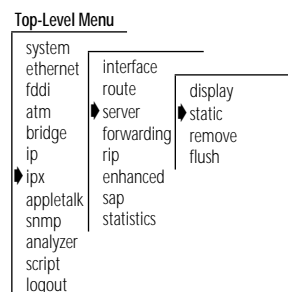
IPX forwarding is enabled, RIP is active, SAP is active

| Interface | Name    | Type | Network  | Node              | Socket | Hops | Age |
|-----------|---------|------|----------|-------------------|--------|------|-----|
| 2         | GB201   | 39b  | 8c141bfe | 08-00-02-04-80-b6 | 8059   | 4    | 73  |
| 2         | GB3COM2 | 39b  | af0bc60f | 00-00-00-00-00-01 | 85fa   | 4    | 85  |

## Defining a Static Server

Before you define static servers on the system, you must define at least one IPX interface. See the section “Defining an IPX Interface” on page 11-3. Static servers remain in the table until you remove them or until you remove the corresponding interface. Static servers take precedence over dynamically learned servers to the same destination. You can have a maximum of eight static servers.

To define a static server:



- 1 From the Administration Console top-level menu, enter:

**ipx server static**

- 2 Enter the interface number of the server.
- 3 Enter the service type of the server.
- 4 Enter the service name of the server.
- 5 Enter the IPX network address of the server.
- 6 Enter the socket value of the server.
- 7 Enter the node address of the server.

- 8 Enter the number of hops to the server.


Example:

```
Enter Interface number: 1
Enter service type: 4
Enter service name: gb201
Enter IPX address: 0x8c14a238
Enter socket: 0x8059
Enter node address: 00-00-2e-f3-56-01
Enter hops: 2
```

### Removing a Server

To remove a server:

#### Top-Level Menu



- 1 From the Administration Console top-level menu, enter:

**ipx server remove**

- 2 Enter the service type of the server.

- 3 Enter the service name of the server.

The server is immediately deleted from the server table.

### Flushing Servers

Flushing deletes all dynamically learned servers from the server table.


To flush all learned servers:

From the Administration Console top-level menu, enter:

**ipx server flush**

All learned servers are immediately deleted from the server table.

#### Top-Level Menu



## Setting IPX Forwarding

You can control whether the system forwards or discards IPX packets addressed to other routers. When you enable IPX forwarding, the system acts as a normal IPX router, forwarding IPX packets from one network to another when required. When you disable IPX forwarding, the system discards any IPX packets not addressed directly to one of its defined IPX interfaces.

### *IPX forwarding default*

By default, IPX forwarding is disabled.

#### Top-Level Menu

```
system
ethernet
fddi
atm
bridge
ip
ipx
appletalk
snmp
analyzer
script
logout
  interface
  route
  server
  forwarding
  rip
  enhanced
  sap
  statistics
```

To enable or disable IPX forwarding:

- 1 From the Administration Console top-level menu, enter:  
**ipx forwarding**
- 2 Enter the IPX forwarding state (**enabled** or **disabled**). To use the value in brackets, press [Return] at the prompt.

## Setting the RIP Mode

You can select a RIP mode that is appropriate for your network. RIP can operate in any of three modes:

- **Off** — The system ignores all incoming RIP packets and does not generate any RIP packets of its own.
- **Passive** — The system processes all incoming RIP packets, but does not broadcast periodic or triggered RIP updates, or respond to RIP requests.
- **Active** — The system processes all incoming RIP packets, responds to explicit requests for routing information, and broadcasts periodic and triggered RIP updates.

*RIP default mode* By default, RIP is *off*.

To set the RIP operating mode:

- 1 From the Administration Console top-level menu, enter:  
**ipx rip**
- 2 Enter the RIP mode (**off**, **passive**, or **active**). To use the value in brackets, press [Return] at the prompt.

#### Top-Level Menu

|           |            |
|-----------|------------|
| system    | interface  |
| ethernet  | route      |
| fdi       | server     |
| atm       | forwarding |
| bridge    | rip        |
| ip        | enhanced   |
| ipx       | sap        |
| appletalk | statistics |
| snmp      |            |
| analyzer  |            |
| script    |            |
| logout    |            |

## Setting the Enhanced RIP Mode

Standard IPX RIP packets can include up to 50 route advertisements, but some routers allow up to 68. Enhanced RIP mode increases the number of entries in a RIP packet that the system will accept. Enhanced RIP mode allows the system greater interoperability with routers that do not explicitly follow the IPX router implementation guidelines.

*Enhanced RIP default* By default, enhanced RIP is *disabled*.

To enable or disable enhanced RIP mode:

- 1 From the Administration Console top-level menu, enter:  
**ipx enhanced**
- 2 Enter the enhanced RIP state (**enabled** or **disabled**). To use the value in brackets, press [Return] at the prompt.

#### Top-Level Menu

|           |            |
|-----------|------------|
| system    | interface  |
| ethernet  | route      |
| fdi       | server     |
| atm       | forwarding |
| bridge    | rip        |
| ip        | enhanced   |
| ipx       | sap        |
| appletalk | statistics |
| snmp      |            |
| analyzer  |            |
| script    |            |
| logout    |            |

## Setting the SAP Mode

You can select a SAP mode that is appropriate for your network. SAP can operate in any of three modes:

- **Off** — The system ignores all incoming SAP packets and does not generate any SAP packets of its own.
- **Passive** — The system processes all incoming SAP packets, but it does *not* broadcast periodic or triggered SAP updates or respond to SAP requests.
- **Active** — The system processes all incoming SAP packets, responds to explicit requests for routing information, and broadcasts periodic and triggered SAP updates.

*SAP default mode* By default, SAP is *off*.

To set the SAP operating mode:

- 1 From the Administration Console top-level menu, enter:  
**ipx sap**
- 2 Enter the SAP mode (**off**, **passive**, or **active**). To use the value in brackets, press [Return] at the prompt.

### Top-Level Menu

|           |            |
|-----------|------------|
| system    | interface  |
| ethernet  | route      |
| fddi      | server     |
| atm       | forwarding |
| bridge    | rip        |
| ip        | enhanced   |
| ipx       | sap        |
| appletalk | statistics |
| snmp      |            |
| analyzer  |            |
| script    |            |
| logout    |            |

## Displaying Statistics

The Administration Console allows you to display four types of IPX-related statistics:


- IPX summary statistics
- IPX RIP statistics
- IPX SAP statistics
- IPX forwarding statistics

### Displaying IPX Summary Statistics

To display IPX summary statistics, from the Administration Console top-level menu, enter:

**ipx statistics summary**

#### Top-Level Menu



Example:

IPX forwarding is enabled, RIP is active, SAP is active

|          |             |         |                |
|----------|-------------|---------|----------------|
| Received | Transmitted | Dropped | Msg Pool Empty |
| 1170878  | 565099      | 0       | 0              |

Table 11-1 describes the IPX summary statistics.

**Table 11-1** IPX Summary Statistics

| Field          | Description  |
|----------------|--|
| Received       | Number of IPX packets received   |
| Transmitted    | Number of IPX packets transmitted  |
| Dropped        | Number of IPX packets dropped  |
| Msg Pool Empty | Number of IPX RIP or IPX SAP messages delivered to the IPX application that were dropped due to resource limitations |



## Displaying IPX RIP Statistics

To display IPX RIP statistics, from the Administration Console top-level menu, enter:

**ipx statistics rip**

Example below:

### Top-Level Menu

```

system
ethernet
fddi
atm
bridge
ip
ipx
appletalk
snmp
analyzer
script
logout
  
```

```

interface
route
server
forwarding
rip
enhanced
sap
statistics
  
```

```

summary
rip
sap
forwarding
  
```

IPX forwarding is enabled, RIP is active, SAP is active

|               |                 |             |
|---------------|-----------------|-------------|
| RIP Received  | RIP Transmitted | RIP dropped |
| 106195        | 7929            | 0           |
| RIP Responses | RIP Requests    | RIP Entries |
| 100552        | 5643            | 2           |

Table 11-2 describes the IPX RIP statistics.

**Table 11-2** IPX RIP Statistics

| Field           | Description  |
|-----------------|--|
| RIP Received    | Number of IPX RIP packets received                   |
| RIP Transmitted | Number of IPX RIP packets transmitted                |
| RIP Dropped     | Number of IPX RIP packets dropped                    |
| RIP Responses   | Number of IPX RIP responses that have been processed |
| RIP Requests    | Number of IPX RIP requests that have been processed  |
| RIP Entries     | Number of routes in the routing table                |


## Displaying IPX SAP Statistics

To display IPX SAP statistics, from the Administration Console top-level menu, enter:

```
ipx statistics sap
```

Example:

### Top-Level Menu



IPX forwarding is enabled, RIP is active, SAP is active

```

SAP Received      SAP Transmitted      SAP dropped
      1064015              22493              0
  
```

```

SAP Responses      SAP Requests      SAP Entries
      1063532              45              0
  
```

```

SAP GNS Responses      SAP GNS Requests
              0              438
  
```

Table 11-1 describes the IPX SAP statistics.

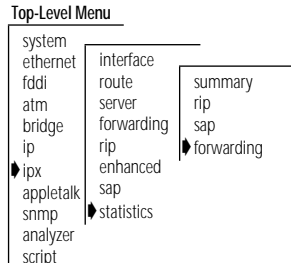
**Table 11-3** IPX SAP Statistics

| Field             | Description   |
|-------------------|---|
| SAP Received      | Number of IPX SAP packets received                                      |
| SAP Transmitted   | Number of IPX SAP packets transmitted                                   |
| SAP Dropped       | Number of IPX SAP packets dropped                                       |
| SAP Responses     | Number of IPX SAP Responses that have been processed                    |
| SAP Requests      | Number of IPX SAP Requests that have been processed                     |
| SAP Entries       | Number of servers in the server table                                   |
| SAP GNS Responses | Number of IPX SAP Get Nearest Service Responses that have been received |
| SAP GNS Requests  | Number of IPX SAP Get Nearest Service Requests processed                |

## Displaying IPX Forwarding Statistics

To display IPX Forwarding statistics, from the Administration Console top-level menu, enter:

**ipx statistics forwarding**



Example:

IPX forwarding is enabled, RIP is active, SAP is active

|                      |                       |                       |
|----------------------|-----------------------|-----------------------|
| Received<br>1335653  | Transmitted<br>565105 | Forwarded<br>0        |
| Hdr Errors<br>13758  | Hop Count Errors<br>0 | Addr Errors<br>13758  |
| No Routes<br>2       | Misc Errors<br>411    |                       |
| NetBIOS Rx<br>150604 | NetBIOS Tx<br>125781  | NetBIOS Max Hops<br>0 |
| Host Rx<br>1171190   | Host Tx<br>565105     |                       |

Table 11-4 describes the IPX forwarding statistics.

**Table 11-4** IPX Forwarding Statistics

| Field            | Description  |
|------------------|--|
| Received         | Number of IPX forwarding packets received                                      |
| Transmitted      | Number of IPX forwarding packets transmitted                                   |
| Forwarded        | Number of IPX packets forwarded by the IPX router                              |
| Hdr Errors       | Number of IPX packets dropped due to IPX Network layer header errors           |
| Hop Count Errors | Number of IPX packets dropped due to exceeded maximum transport control        |
| Addr Errors      | Number of IPX packet dropped due to IPX Address errors in network layer header |
| No Routes        | Number of IPX packets dropped because the IPX route is unknown                 |
| Misc Errors      | Number of multicasts attempted to be forwarded                                 |
| NetBIOS Rx       | Number of IPX NetBIOS packets received   |
| NetBIOS Tx       | Number of IPX NetBIOS packets transmitted                                      |
| NetBIOS Max Hops | Number of IPX NetBIOS packets that exceeded the transport control maximum      |
| Host Rx          | Number of IPX packets delivered to the IPX host's RIP and SAP applications     |
| Host Tx          | Number of IPX packets transmitted from IPX host's RIP and SAP applications     |

# 12

## ADMINISTERING APPLE TALK® ROUTING

This chapter describes how to set up your LANplex® system to use the AppleTalk protocol to route packets. For more information on how AppleTalk routing works, see Chapter 7: Routing with AppleTalk.

You can display and configure the following:

- AppleTalk interfaces
- Routes
- AARP cache
- Zones
- AppleTalk forwarding
- Checksum generation and verification
- AppleTalk statistics

## Administering Interfaces

An AppleTalk interface defines the relationship between an AppleTalk Virtual LAN (VLAN) and the AppleTalk network. Every AppleTalk interface has one AppleTalk VLAN associated with it. Each switching module has one AppleTalk interface defined for each subnet directly connected to it.



*You must first define a VLAN, as described in Chapter 8, before you define an associated AppleTalk interface.*



*You can configure a maximum of 32 interfaces per router.*

An AppleTalk interface has several elements associated with it:

- **Seed Interface** — You can configure the interface to be a seed interface or non-seed interface. Seed interfaces initialize the network with the configuration information the administrator enters. These include network range, address, zone name, and ports. Non-seed interfaces wait and listen for a seed interface and then take this configuration initialization information from the first seed interface they detect. After the non-seed interface obtains a network configuration, it begins to participate in the routing of the network.
- **Network Range** — A range of numbers used to designate a network segment's identity. This element allows the physical segment between two LANplex systems to be a range of multiple networks.
- **Address** — The AARP address based on the network range and the network node (1-253).
- **Zone** — The default zone name, as well as up to 15 additional defined zones.
- **State** — This is the status of the AppleTalk interface, which indicates whether the interface is available for communications (*up*) or unavailable (*down*).
- **VLAN Index** — The number of the AppleTalk VLAN associated with the AppleTalk interface. The VLAN index indicates which bridge ports are associated with the AppleTalk interface. When the menu prompts you for this option, it displays a list of available VLAN indexes and their ports.

## Displaying AppleTalk Interfaces

You can display a table that shows all AppleTalk interfaces and their parameter settings configured for the system.

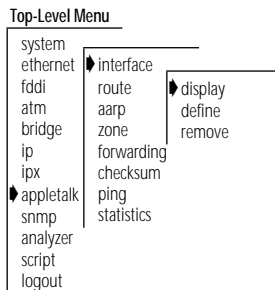
To display the AppleTalk interfaces defined on the router, from the Administration Console top-level menu, enter:

```
appletalk interface display
```

Example:

DDP forwarding is enabled.

| Index | Network Range | Address  | State   | VLAN index |
|-------|---------------|----------|---------|------------|
| 1     | 20112-20112   | 20112.27 | enabled | 3          |
| 2     | 20124-20124   | 20124.1  | enabled | 2          |
| 3     | 20125-20125   | 20125.1  | enabled | 4          |



## Defining an Interface

When you define an interface, you define the interface's network range, zone name, and the VLAN index associated with the interface. You must define an AppleTalk VLAN before you define the AppleTalk interface to associate with that VLAN.

To define an AppleTalk interface:

- 1 At the Administration Console's top-level menu, enter:

```
appletalk interface define
```

You are prompted for the interface's parameters. To use the value in brackets, press [Return] at the prompt.

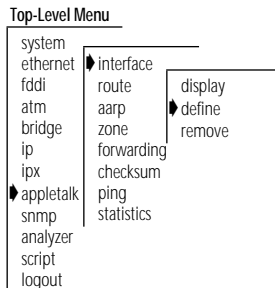
The following message appears:

```
Configure seed interface? (n,y) [y]:
```

- 2 Enter **n** (no) or **y** (yes).
- 3 Enter the start of the network range associated with the interface.
- 4 Enter the end of the network range associated with the interface.
- 5 Enter the default zone name.



*The default zone name is used by clients that have not been configured to use a particular zone.*



- 6 Enter the zone name.



*You can enter up to 16 zone names per interface.*

- 7 Type **q** after entering all the zone names.
- 8 Enter the index of the AppleTalk VLAN associated with this interface.

Example:

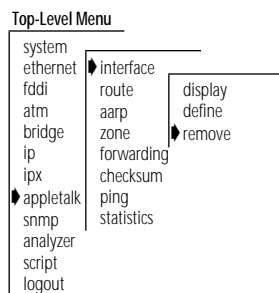
```
Enter Start of Network Range:10000
Enter End of Network Range: 10100
Enter Default Zone:engineering
Enter Zone Name:q
Appletalk VLANs:
      Index      Ports
        3         1-8
        4         9-12
Select VLAN index: 3
```

## Removing an Interface

You might want to remove an interface if you no longer perform routing on the ports associated with the interface.

To remove an AppleTalk interface:

- 1 At the Administration Console's top-level menu, enter:  
**appletalk interface remove**
- 2 Enter the index number(s) of the interface(s) you want to remove.  
The interface is no longer defined on the router.





## Administering Routes

Your system maintains a table of routes to other AppleTalk networks. The routing table is generated automatically by the Routing Table Maintenance Protocol (RTMP). RTMP defines 1) the rules for exchanging information between routers so that the routers can maintain their routing tables, and 2) the rules for the information contained within each routing table.

Each routing table entry contains the following information:

- **Network Range**

A range of numbers used to designate a network segment's identity

- **Distance**

The distance in hops to the destination network

- **Interface**

The defined interface number

- **State**

The status (good, suspect, bad, or really bad) of each route

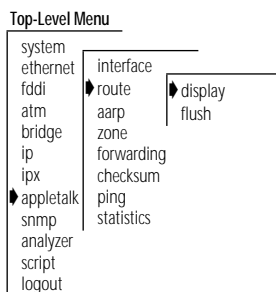
## Displaying the Routing Table

You can display the routing tables for the system to determine which routes are configured and if they are operational.

To display the contents of the routing table:

From the Administration Console top-level menu, enter:

**appletalk route display**




The following example shows a routing table display:

| Network | Range       | Distance | Interface | State |
|---------|-------------|----------|-----------|-------|
|         | 1-1         | 10       | 1         | good  |
|         | 3           | 4        | 1         | good  |
|         | 10-14       | 6        | 1         | good  |
|         | 15-19       | 7        | 1         | good  |
|         | 61          | 6        | 1         | good  |
|         | 100-100     | 10       | 1         | good  |
|         | 201-300     | 7        | 1         | good  |
|         | 2010-2015   | 2        | 1         | good  |
|         | 10009-10009 | 5        | 1         | good  |
|         | 10010-10010 | 7        | 1         | good  |
|         | 10060-10060 | 8        | 1         | good  |
|         | 10110-10113 | 5        | 1         | good  |
|         | 10116-10117 | 5        | 1         | good  |
|         | 10118-10118 | 6        | 1         | good  |
|         | 10119-10119 | 4        | 1         | good  |
|         | 10120-10120 | 7        | 1         | good  |
|         | 10122-10122 | 9        | 1         | good  |
|         | 10310-10329 | 10       | 1         | good  |
|         | 10410-10410 | 8        | 1         | good  |
|         | 11010-11019 | 9        | 1         | good  |

## Flushing all Routes

Flushing deletes all dynamically learned routes from the routing table.

### Top-Level Menu



To flush all learned routes:

- 1 At the Administration Console's top-level menu, enter:

**appletalk route flush**

---

## Administering the AARP Cache

AARP allows hardware addresses to be mapped to an AppleTalk protocol address. AppleTalk uses dynamically assigned 24-bit addresses, unlike the statically-assigned 48-bit addresses used by Ethernet and token ring.

To make the address mapping process easier, AARP uses an Address Mapping Table (AMT). The most recently used addresses are maintained in the AMT. If an address is not in the AMT, AARP sends a request to the desired protocol address and the hardware address is added to the table when the destination node replies.

AARP is also responsible for registering a node's dynamically assigned address on the network. This process is described below:

- AARP randomly assigns an address.
- AARP broadcasts AARP probe packets to this address to determine if another node is already using the address.
- If there is no reply, the address becomes that node's address.
- If there is a reply, AARP repeats this process until an available address is discovered.

In the Administration Console, you can:

- Display the cache
- Remove entries
- Flush the cache

## Displaying the AARP Cache

You can display the AARP cache for the system to determine which routes are configured and if they are operational.


To display the contents of the AARP cache:

From the Administration Console top-level menu, enter:

**appletalk aarp display**

The following example shows an AARP cache display:

### Top-Level Menu

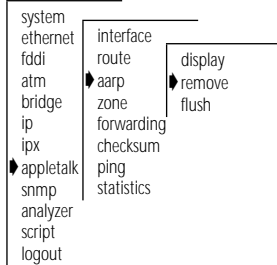


| AARP Address | MAC Address       | Interface | Age (secs) |
|--------------|-------------------|-----------|------------|
| 20112.125    | 00-20-af-0b-e1-7c | 1         | 211        |
| 20112.177    | 00-00-89-01-91-f0 | 1         | 20         |
| 20112.192    | 00-00-89-01-91-f3 | 1         | 6          |
| 20112.150    | 00-00-89-01-8b-51 | 1         | 18         |
| 20112.1      | 08-00-02-04-80-b6 | 1         | 31         |
| 20125.193    | 08-00-07-d7-69-1f | 3         | 388        |
| 20125.76     | 08-00-07-66-62-9d | 3         | 862        |
| 20125.67     | 08-00-07-ee-10-a2 | 3         | 851        |
| 20124.41     | 08-00-07-7c-c3-d8 | 2         | 864        |
| 20112.225    | 00-20-af-0b-d8-f1 | 1         | 270        |
| 20112.135    | 00-20-af-9e-68-62 | 1         | 174        |
| 20112.147    | 00-00-94-41-de-79 | 1         | 26         |
| 20112.132    | 08-00-09-7f-98-c5 | 1         | 24         |
| 20112.112    | 08-00-07-7c-20-61 | 1         | 121        |
| 20112.148    | 08-00-07-ac-56-4b | 1         | 1098       |
| 20112.244    | 00-20-af-0b-ff-72 | 1         | 35         |
| 20112.21     | 08-00-07-dc-e5-c4 | 1         | 8932       |
| 20112.131    | 08-00-07-54-88-b1 | 1         | 397        |
| 20124.35     | 08-00-07-57-ec-58 | 2         | 368        |
| 20112.97     | 08-00-07-9e-09-86 | 1         | 1925       |
| 20112.4      | 08-00-07-ec-98-3d | 1         | 121        |
| 20112.180    | 08-00-07-f7-cf-de | 1         | 110        |
| 20112.108    | 08-00-07-4f-74-7e | 1         | 5833       |
| 20112.56     | 08-00-07-bc-10-fc | 1         | 120        |
| 20112.110    | 00-40-10-56-1a-b5 | 1         | 110        |
| 20112.155    | 08-00-07-6c-88-77 | 1         | 5536       |
| 20112.243    | 08-00-07-66-72-c7 | 1         | 4940       |
| 20112.253    | 08-00-20-12-75-bf | 1         | 70         |
| 20125.104    | 08-00-07-66-2b-c2 | 3         | 848        |
| 20112.236    | 00-80-3e-02-81-66 | 1         | 3841       |

## Removing an Entry in the Cache

To remove an AARP cache entry:

### Top-Level Menu



- 1 At the Administration Console's top-level menu, enter:

**appletalk aarp remove**


- 2 Enter the AARP address at the prompt.

The entry is removed.

## Flushing All Cache Entries

To flush all AARP cache entries:

### Top-Level Menu



- 1 At the Administration Console's top-level menu, enter:

**appletalk aarp flush**

## Displaying the Zone Table

AppleTalk allows for the logical grouping of nodes into zones to make navigation through the network easier. This is done with the Zone Information Protocol (ZIP). ZIP helps routers maintain a mapping of network numbers to zones in the entire network. To do this, ZIP creates and maintains a Zone Information Table (ZIT) in each router. The entries in this table match the network numbers with the zone names.

In the Administration Console, you can display the zone table either by network numbers or by zones.

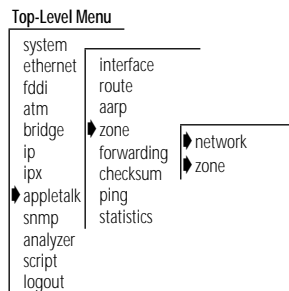
To display the zone table:

From the Administration Console top-level menu, enter:

**appletalk zone display network**

OR

**appletalk zone display zone**



Depending on the command entered, the zone table is displayed by network or zone. An example of each type of display is shown below:

Zone Table by Network Numbers

```
DDP forwarding is enabled.

Network 1-1 has 1 known zone
  Munich GmbH

Network 3 has 1 known zone
  Ethernet A5D85800

Network 10-14 has 1 known zone
  Freds_Ethernet

Network 15-19 has 1 known zone
  Freds-Token

Network 61 has 1 known zone
  DevMacNet

Network 100-100 has 1 known zone
  France Les Ulis

Network 201-300 has 1 known zone
  Fred_Wilma

Network 2010-2015 has 1 known zone
  NY

Network 10009-10009 has 2 known zones
  Hemel NSOPS
  3Com Arpeggio

Network 10010-10010 has 1 known zone
  Marlow EUR
```

Zone Table by Zones

```
DDP forwarding is enabled.

Zone Holmdel is assigned to 2 networks
  21105-21105
  21010-21010

Zone NY is assigned to 2 networks
  63535-63535
  2010-2015

Zone Manchester UK is assigned to 1 network
  10310-10329

Zone DC8 is assigned to 1 network
  30110-30129

Zone Chicago is assigned to 1 network
  22030-22030

Zone Startek-Enet1 is assigned to 1 network
  20033-20033

Zone Startek-TR1 is assigned to 1 network
  20037-20037

Zone Test GmbH is assigned to 1 network
  12010-12012

Zone Madrid3Com is assigned to 1 network
  14010-14029

Zone NSDEng is assigned to 1 network
  32910-32910
```

Configuring Forwarding

Top-Level Menu

system  
ethernet  
fddi  
atm  
bridge  
ip  
ipx  
♦ appletalk  
snmp  
analyzer  
script  
logout

interface  
route  
aarp  
zone  
♦ forwarding  
checksum  
ping  
statistics

You can control whether the router forwards or discards AppleTalk packets addressed to other hosts. When you enable forwarding, the router processes packets as usual, forwarding AppleTalk packets from one subnet to another when required. When you disable IP forwarding, the router discards any AppleTalk packets not addressed directly to one of its defined interfaces.

- 1 At the Administration Console's top-level menu, enter:  
**appletalk forwarding**
- 2 Enter **enable** or **disable** at the prompt.

## Configuring Checksum

Checksum is a simple method used for detecting errors in the transmission of data. Checksum generation totals the bytes comprising the data and adds this sum to the end of the data packet. Checksum verification allows you to verify the integrity of the data that is routed. You can enable or disable checksum generation and verification states.

To enable or disable checksum generation/verification:

- 1 At the Administration Console's top-level menu, enter:  
**appletalk checksum**
- 2 Enter **enable** or **disable** at the checksum generation prompt.
- 3 Enter **enable** or **disable** at the checksum verification prompt.

### Top-Level Menu

|             |            |
|-------------|------------|
| system      | interface  |
| ethernet    | route      |
| fdi         | aarp       |
| atm         | zone       |
| bridge      | forwarding |
| ip          | checksum   |
| ipx         | ping       |
| ♦ appletalk | statistics |
| snmp        |            |
| analyzer    |            |
| script      |            |
| logout      |            |

## Pinging an AppleTalk Node

The AppleTalk Echo Protocol (AEP) sends a datagram (an Echo Request) from one node to another, which causes the destination node to return or *echo*, the datagram (an Echo Reply) to the sender. This allows you to determine whether a node is accessible before any sessions are started.

To ping an AppleTalk node:

- 1 At the Administration Console's top-level menu, enter:  
**appletalk ping**  
You are prompted for a node address.
- 2 Enter the address of the node you want to ping.  
If the node is accessible, you receive a response.

### Top-Level Menu

|             |            |
|-------------|------------|
| system      | interface  |
| ethernet    | route      |
| fdi         | aarp       |
| atm         | zone       |
| bridge      | forwarding |
| ip          | checksum   |
| ipx         | ping       |
| ♦ appletalk | statistics |
| snmp        |            |
| analyzer    |            |
| script      |            |
| logout      |            |

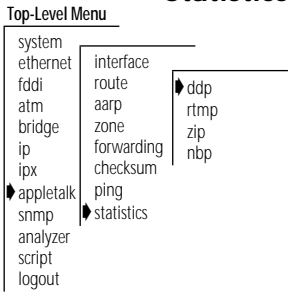


## Viewing Appletalk Statistics

You can view statistics specific to the following AppleTalk protocols:

- Datagram Delivery Protocol (DDP)
- Routing Table Maintenance Protocol (RTMP)
- Zone Information Protocol (ZIP)
- Name Binding Protocol (NBP)

### Displaying DDP Statistics



To display DDP statistics:

From the Administration Console top-level menu, enter:

**appletalk statistics ddp**

The following is an example of DDP summary statistics:

```
DDP forwarding is enabled.

      inReceives      inForwards      inLocals      inNoRoutes
          131131          113171          17906           22

      inNoClients      inTooShorts      inTooLongs      inShortDdps
           0              0              0              0

      inCsumErrors      inBcastErrors      inTooFars      inDiscards
           0              0              0             54

      outLocals
          15600
```

Table 12-1 describes the AppleTalk DDP statistics you can view.

**Table 12-1** AppleTalk Statistics

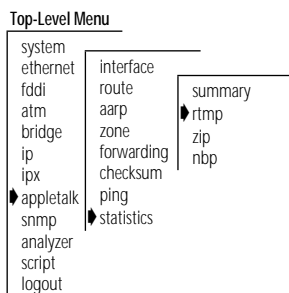
| Field       | Description   |
|-------------|---|
| inReceives  | Total number of packets received, including those with errors                 |
| inForwards  | Total number of packets forwarded, including those with errors                |
| inLocals    | Number of DDP datagrams for which this entity was their final DDP destination |
| inNoRoutes  | Number of DDP datagrams dropped because a route could not be found            |
| inNoClients | Number of DDP datagrams dropped because of an unknown DDP type                |

continued

**Table 12-1** AppleTalk Statistics (continued)

| Field         | Description  |
|---------------|--|
| inTooShorts   | Number of input DDP datagrams dropped because the received data length was less than the data length specified in the DDP header or the received data length was less than the length of the expected DDP header |
| inTooLongs    | Number of input DDP datagrams dropped because they exceeded the maximum DDP datagram size  |
| inShortDdps   | Number of input DDP datagrams dropped because this entity was not their final destination and their type was short DDP   |
| inCsumErrors  | Number of DDP datagrams which were dropped because of a checksum error   |
| inBcastErrors | Number of DDP datagrams for which this DDP entity was their final destination, and which were dropped because of a broadcast error   |
| inTooFars     | Number of input datagrams dropped because this entity was not their final destination and their hop count would exceed 15  |
| inDiscards    | Number of DDP Datagrams thrown out during the routing process  |
| outLocals     | Number of host-generated DDP datagrams   |

## Displaying RTMP Information



To display RTMP statistics:

From the Administration Console top-level menu, enter:

**appletalk statistics rtmp**

An example of summary statistics is shown below:

```

DDP forwarding is enabled.

      inDatas      inRequests      outDatas      outRequests
        7204             0         4865             6

routeEqChgs routeLessChgs routeDeletes routeOverflows
         0             0             0             0

inVersionErrors inOtherErrors
         0             119
  
```

Table 12-2 describes the RTMP statistics you can view.

**Table 12-2** RTMP Statistics

| Field          | Description   |
|----------------|---|
| inDats         | Number of good RTMP data packets received   |
| inRequests     | Number of good RTMP request packets received  |
| outDats        | Number of good RTMP data packets sent   |
| outRequests    | Number of RTMP request packets sent   |
| routeEqChgs    | Number of times RTMP changes the Next Internet Router in a routing entry because the hop count advertised in a routing table was equal to the current hop count for a particular network  |
| routeLessChgs  | Number of times RTMP changes the Next Internet Router in a routing entry because the hop count advertised in a routing table was less than the current hop count for a particular network |
| routeDeletes   | Number of times RTMP deletes a route because it was aged out of the table   |
| routeOverflows | Number of times RTMP attempted to add a route to the RTMP table but failed due to lack of space   |
| inVersionErrs  | Number of RTMP packets received that were rejected due to a version mismatch  |
| inOtherErrs    | Number of RTMP packets received that were rejected for an error other than due to a version mismatch  |


## Displaying ZIP Information

To display ZIP statistics:

From the Administration Console top-level menu, enter:

**appletalk statistics zip**

### Top-Level Menu



An example of summary statistics is shown below:

DDP forwarding is enabled.

|               |               |              |                |
|---------------|---------------|--------------|----------------|
| inQueries     | inReplies     | inExReplies  | inGniRequests  |
| 248           | 14            | 0            | 182            |
| inGniReplies  | inLocalZones  | inZoneLists  |                |
| 22            | 30            | 0            |                |
| inObsoletes   | inZoneCons    | inZoneInvs   | inErrors       |
| 0             | 0             | 22           | 0              |
| outQueries    | outReplies    | outExReplies | outGniRequests |
| 9             | 0             | 277          | 13             |
| outGniReplies | outLocalZones | outZoneLists |                |
| 182           | 0             | 30           |                |
| outZoneInvs   | outAddrInvs   |              |                |

Table 12-3 describes the ZIP statistics you can view:

**Table 12-3** ZIP Statistics

| Field          | Description   |
|----------------|---|
| inQueries      | Number of ZIP queries received  |
| inReplies      | Number of ZIP replies received  |
| inExReplies    | Number of ZIP extended replies received   |
| inGniRequests  | Number of ZIP GetNetInfo request packets received   |
| inGniReplies   | Number of ZIP GetNetInfo reply packets received   |
| inLocalZones   | Number of Zip GetLocalZones requests packets received   |
| inZoneLists    | Number of Zip GetZoneLists requests packets received  |
| inObsoletes    | Number of ZIP Takedown or ZIP Bringup packets received  |
| inZoneCons     | Number of times a conflict has been detected between this entity's zone information and another entity's zone information   |
| inZoneInvs     | Number of times this entity has received a ZIP GetNetInfo reply with the zone invalid bit set because the corresponding GetNetInfo request had an invalid zone name |
| inErrors       | Number of ZIP packets received that were rejected for any error   |
| outQueries     | Number of ZIP queries sent  |
| outReplies     | Number of ZIP replies sent  |
| outExReplies   | Number of ZIP extended replies sent   |
| outGniRequests | Number of ZIP GetNetInfo packets sent   |

continued

Table 12-3 ZIP Statistics (continued)

| Field         | Description   |
|---------------|---|
| outGniReplies | Number of ZIP GetNetInfo reply packets sent out of this port  |
| outzoneInvs   | Number of times this entity has sent a ZIP GetNetInfo reply with the zone invalid bit set in response to a GetNetInfo request with an invalid zone name |
| outAddrInvs   | Number of times this entity had to broadcast a ZIP GetNetInfo reply because the GetNetInfo request had an invalid address                               |

Displaying NBP Information

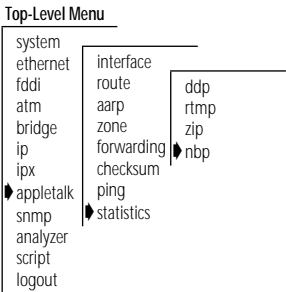
The NBP handles the translations between the numeric internet address and the alphanumeric entity names used by AppleTalk.

To display NBP statistics:

From the Administration Console top-level menu, enter:

**appletalk statistics nbp**

An example of summary statistics is shown below:



```
> forwarding is enabled.


      inLkupReqs      inBcastReqs      inFwdReqs      inLkupReplies
          3093              611          5951              0

      inErrors
          0
```

Table 12-4 describes the NBP statistics you can view.


**Table 12-4** NBP Statistics

| Field         | Description   |
|---------------|---|
| inLkupReqs    | Number of NBP Lookup Requests received                          |
| inBcastsReqs  | Number of NBP Broadcast Requests received                       |
| inFwdReqs     | Number of NBP Forward Requests received                         |
| inLkupReplies | Number of NBP Lookup Replies received                           |
| inErrors      | Number of NBP packets received that were rejected for any error |



# REMOTE MONITORING (RMON) AND THE LANPLEX® SYSTEM

**Chapter 13** Remote Monitoring (RMON) Technology



# 13

## REMOTE MONITORING (RMON) TECHNOLOGY

This chapter provides an overview of RMON and describes the specific LANplex® RMON implementation.

---

### What Is RMON?

The Remote Monitoring (RMON) Management Information Base (MIB) provides a way to monitor and analyze a local area network LAN from a remote location. RMON is defined by the Internet Engineering Task Force (IETF) in documents RFC 1271 and RFC 1757. A typical RMON implementation has two components:

- **The Probe** — Connects to a LAN segment, examines all the LAN traffic on that segment and keeps a summary of statistics (including historical data) in its local memory.
- **The Management Console** — Communicates with the probe and collects the summarized data from it. The console does not need to reside on the same network as the probe, and can manage the probe through SNMP.

The RMON specification consists almost entirely of the definition of the MIB. The RMON MIB contains standard MIB variables defined to collect comprehensive network statistics that alert a network administrator to significant network events. If the embedded RMON agent operates full time, it will collect data on the correct port at the time the relevant network event occurs.

This chapter includes the following information about RMON:

- Benefits of RMON
- LANplex RMON implementation
- The Management Information Base (MIB)
- Alarms



---

## Benefits of RMON

Traditional network management applications poll network devices such as switches, bridges, and routers at regular intervals from a network management console. The console gathers statistics, identifies trends, and can highlight network events. The console polls network devices constantly to determine if the network is within its normal operating conditions.

As network size and traffic levels grow, however, the network management console can become overburdened by the amount of data it must collect. Frequent console polling also generates significant network traffic that itself can create problems for the network itself.

An RMON implementation offers solutions to both of these problems:

- The RMON probe looks at the network on behalf of the network management console without affecting the characteristics and performance of the network itself.
- The RMON MIB reports by exception rather than by sending constant or frequent information to the network management console. The RMON probe informs the network management console directly if the network enters an abnormal state. The console can then use more information from the probe, such as history information, to diagnose the abnormal condition.

---

## LANplex RMON Implementation

The LANplex Extended Switching software offers full time embedded RMON support through SNMP for four RMON Groups. When combined with the Roving Analysis Port (RAP) function, RMON support for these groups provides a comprehensive and powerful mechanism for managing your network.



*You can gain access to the RMON capabilities of the LANplex 2500 system only through SNMP applications such as Transcend® Enterprise Manager software, not through the serial interface or telnet. For more information about the details of managing 3Com devices using RMON, see the user documentation of the Transcend Network Management Application for Windows.*

The LANplex system supports four of the RMON groups defined by the IETF. Table 13-1 lists these supported groups.

**Table 13-1** RMON Groups Supported in the LANplex® System

| Group      | Group Number | Purpose   |
|------------|--------------|---|
| Statistics | 1            | Maintains utilization and error statistics for the segment being monitored                    |
| History    | 2            | Gathers and stores periodic statistical samples from the statistics group.                    |
| Alarm      | 3            | Allows you to define thresholds for any MIB variable and trigger an alarm.                    |
| Events     | 9            | Allows you to define actions based on alarms. You can generate traps, log the alarm, or both. |

### 3Com Transcend RMON Agents

RMON requires one probe per LAN segment. Because a segment is a portion of the LAN separated by a bridge or router, the cost of implementing many probes in a large network can be high.

To solve this problem, 3Com has built an inexpensive RMON probe into the Transcend SmartAgent software in each LANplex 2500 system. This probe allows you to deploy RMON widely around the network at a cost no more than that for traditional network monitors.

Placing probe functionality inside the LANplex 2500 system has these advantages:

- You can integrate RMON with normal device management
- The LANplex system can manage conditions proactively

The LANplex system associates statistics with individual ports and then takes action based on these statistics. For example, the system can generate a log event and send an RMON trap if errors on a port exceed a user-set threshold.



*You must assign an IP address to the LANplex system to manage RMON. See the LANplex® 2500 Administration Console User Guide for information on how to assign an IP address.*

Figure 13-1 shows an example of a LANplex RMON implementation. The LANplex 2500 system in this figure has two Fast Ethernet connections in addition to the 10BASE-T connections.

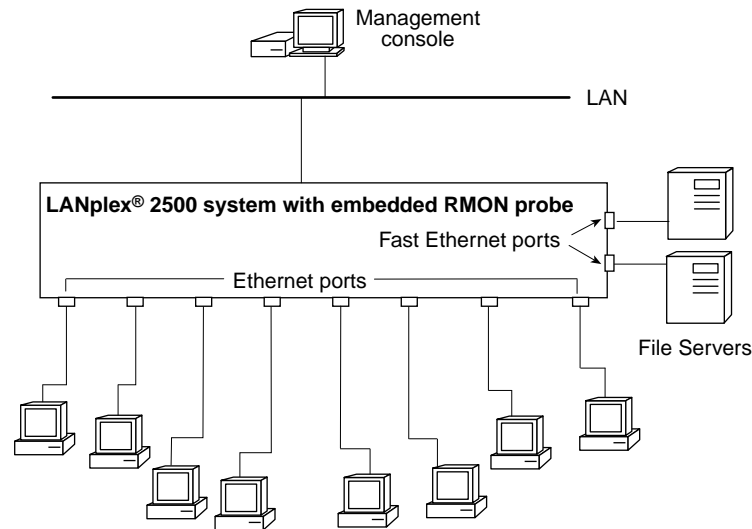


Figure 13-1 Embedded RMON Implemented on the LANplex System

## Management Information Base (MIB)

A MIB is a structured set of data that describes the way the network is functioning. The management software, known as the *agent*, gains access to this set of data and extracts the information it needs. The agent can also store data in the MIB.

The organization of a MIB allows a Simple Network Management Protocol (SNMP) network management package such as the Transcend Enterprise Manager application suite to manage a network device without a specific description of that device. 3Com ships SNMP MIB files with LANplex Extended Switching System software as ASN.1 files.

### MIB Objects

The data in the MIB consists of objects that represent features of the equipment that an agent can control and manage. Examples of objects in the MIB include a port that you can enable or disable and a counter that you can read.

A counter is a common type of MIB object used by RMON. A counter object might record the number of frames transmitted onto the network. The MIB might contain an entry for the counter object something like the one in Figure 13-2 for the counter object.

```

etherStatsPkts OBJECT-TYPE
    SYNTAX      Counter
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        This is a total number of packets
        received, including bad packets,
        broadcast packets, and multicast
        packets.
    ::= { etherStatsEntry 5 }

```

**Figure 13-2** Example of an RMON MIB Counter Object

The displayed information includes these items:

- The formal name of the counter is *etherStatsPkts*. (Ethernet, Statistics, Packets)
- The access is read-only
- The number of the counter's column in the table: 5

The name of the table in which the counter resides is *3CometherStatTable*, although this does not appear in the display.

You do not need to know the contents of every MIB object to manage a network. Most network management applications, including Transcend Enterprise Manager Software, make the MIB transparent. However, knowing how different management features are derived from the MIB allows you to better understand how to use the information that they provide.

---

## Alarms

The LANplex system supports the following syntax for alarms: counters, gauges, integers and timeticks. These mechanisms report information about the network to the network administrator. Counters, for example, hold and update the number of occurrences of a particular event through a port, module, or switch on the network. Alarms monitor the counters and report instances of when counters exceed their set threshold.

Counters are useful when you compare their values at specific time intervals to determine rates of change. The time intervals can be short or long, depending on what you measure. Occasionally, reading counters can give you misleading results.

Counters are not infinite, which makes rate comparisons an efficient way to use them. When counters reach a predetermined limit, they return to 0 (*roll over*). A single low counter value might accurately represent a condition on the network. It might simply indicate that a roll over has occurred.



*When you disable a port, the application might not update some of the statistics counters associated with it.*

An alarm calculates the difference in counter values over a set time interval and remembers the high and low values. When the value of a counter exceeds a preset threshold, the alarm reports this occurrence.

You can assign alarms with Transcend Enterprise Manager or any other SNMP network management application to monitor any counter, gauge, time tick, or integer. Consult the documentation for your management application for details on setting up alarms.

### Setting Alarm Thresholds

Thresholds determine when an alarm reports that a counter has exceeded a certain value. You can set alarm thresholds through the network manually, and choose any value for them that is appropriate for your application. The network management software monitors the counters and thresholds continually during normal operations to provide data for later calibration.

### Example of an Alarm Threshold

Figure 13-3 shows a counter with thresholds set manually.




Figure 13-3 Manually Set Thresholds


You can associate an alarm with the high threshold, the low threshold, or both. The actions taken because of an alarm depend on the network management application.

**RMON Hysteresis Mechanism**

The RMON hysteresis mechanism provides a way to prevent small fluctuations in counter values from causing alarms. This mechanism generates an alarm only under the following conditions:


- The counter value exceeds the high threshold after previously exceeding the low threshold. (An alarm does not occur if the value has not fallen below the low threshold before rising above the high threshold.)
- The counter value exceeds the low threshold after previously exceeding the high threshold. (An alarm does not occur if the value has not risen above the high threshold before falling below the low threshold.)

In Figure 13-3, for example, an alarm occurs the first time the counter exceeds the high threshold, but not at the second time. At the first instance, the counter is rising from the low threshold, while in the second instance, it is not.




# APPENDIX

## Appendix A Technical Support







## TECHNICAL SUPPORT

3Com provides easy access to technical support information through a variety of services. This appendix describes these services.

---

### On-line Technical Services

3Com offers worldwide product support 24 hours a day, seven days a week, through the following on-line systems:

- 3Com Bulletin Board Service (3ComBBS)
- World Wide Web site
- 3ComForum on CompuServe® online service
- 3ComFacts<sup>SM</sup> automated fax service

### 3Com Bulletin Board Service

3ComBBS contains patches, software, and drivers for all 3Com products, as well as technical articles. This service is available via analog modem or ISDN 24 hours a day, seven days a week.

#### Access by Analog Modem

To reach the service by modem, set your modem to 8 data bits, no parity, and 1 stop bit. Call the telephone number nearest you:

| Country              | Data Rate       | Telephone Number                                    |
|----------------------|-----------------|---|
| Australia            | up to 14400 bps | (61) (2) 9955 2073                                  |
| France               | up to 14400 bps | (33) (1) 69 86 69 54                                |
| Germany              | up to 9600 bps  | (49) (89) 627 32 188 <b>or</b> (49) (89) 627 32 189 |
| Hong Kong            | up to 14400 bps | (852) 2537 5608                                     |
| Italy (fee required) | up to 14400 bps | (39) (2) 273 00680                                  |
| Japan                | up to 14400 bps | (81) (3) 3345 7266                                  |
| Singapore            | up to 14400 bps | (65) 534 5693                                       |
| Taiwan               | up to 14400 bps | (886) (2) 377 5840                                  |
| U.K.                 | up to 28800 bps | (44) (1442) 278278                                  |
| U.S.                 | up to 28800 bps | (1) (408) 980 8204                                  |

### Access by Digital Modem

ISDN users can dial in to 3ComBBS using a digital modem for fast access up to 56 Kbps. To access 3ComBBS using ISDN, dial the following number:

**(408) 654-2703**

**World Wide Web Site** Access the latest networking information on 3Com's World Wide Web site by entering our URL into your Internet browser:

**`http://www.3Com.com/`**

This service features news and information about 3Com products, customer service and support, 3Com's latest news releases, selected articles from 3TECH™ journal (3Com's award-winning technical journal), and more.

**3ComForum on CompuServe®** 3ComForum is a CompuServe® service containing patches, software, drivers, and technical articles about all 3Com products, as well as a messaging section for peer support. To use 3ComForum, you need a CompuServe account.

To use 3ComForum:

- 1** Log on to CompuServe.
- 2** Enter **go threecom**
- 3** Press [Return] to see the 3ComForum Main menu.

### 3ComFacts™ Automated Fax Service

3Com Corporation's interactive fax service, 3ComFacts, provides data sheets, technical articles, diagrams, and troubleshooting instructions on 3Com products 24 hours a day, seven days a week.

Call 3ComFacts using your Touch-Tone® telephone at one of these international access numbers:

| Country   | Telephone Number   |
|-----------|--------------------|
| Hong Kong | (852) 2537 5610    |
| U.K.      | (44) (1442) 278279 |
| U.S.      | (1) (408) 727 7021 |

Local access numbers are available within the following countries:

| Country   | Telephone Number | Country              | Telephone Number |
|-----------|------------------|----------------------|------------------|
| Australia | 800 123853       | Netherlands          | 06 0228049       |
| Belgium   | 0800 71279       | Norway               | 800 11062        |
| Denmark   | 800 17319        | Portugal             | 0505 442607      |
| Finland   | 98 001 4444      | Russia (Moscow only) | 956 0815         |
| France    | 05 90 81 58      | Spain                | 900 964445       |
| Germany   | 0130 8180 63     | Sweden               | 020 792954       |
| Italy     | 1678 99085       | U.K.                 | 0800 626403      |

### Support from Your Network Supplier

If additional assistance is required, contact your network supplier. Many suppliers are authorized 3Com service partners who are qualified to provide a variety of services, including network planning, installation, hardware maintenance, application training, and support services.

When you contact your network supplier for assistance, have the following information ready:

- Diagnostic error messages
- A list of system hardware and software, including revision levels
- Details about recent configuration changes, if applicable

If you are unable to contact your network supplier, see the following section on how to contact 3Com.

## Support from 3Com

If you are unable to receive support from your network supplier, technical support contracts are available from 3Com.

In the U.S. and Canada, call **(800) 876-3266** for customer service.

If you are outside the U.S. and Canada, contact your local 3Com sales office to find your authorized service provider. Use one of these numbers:

| Country    | Telephone Number   | Country              | Telephone Number   |
|------------|--------------------|----------------------|--------------------|
| Australia* | 1800 678 515       | Japan                | (81) (3) 3345 7251 |
|            |                    | Mexico               | (525) 531 0591     |
| Belgium*   | 0800 71429         | Netherlands*         | 06 0227788         |
| Brazil     | (55) (11) 546 0869 | Norway*              | 800 11376          |
| Canada     | (416) 498 3266     | Singapore            | (65) 538 9368      |
| Denmark*   | 800 17309          | South Africa         | (27) (11) 803 7404 |
| Finland*   | 0800 113153        | Spain*               | 900 983125         |
| France*    | 05 917959          | Sweden*              | 020 795482         |
| Germany*   | 0130 821502        | Taiwan               | (886) (2) 577 4352 |
| Hong Kong  | (852) 2501 1111    | United Arab Emirates | (971) (4) 349049   |
| Ireland*   | 1 800 553117       | U.K.*                | 0800 966197        |
| Italy*     | 1678 79489         | U.S.                 | (1) (408) 492 1790 |

\* These numbers are toll-free.

## Returning Products for Repair

Before you return a product directly to 3Com for repair, you must first call for a Return Materials Authorization (RMA) number. A product sent to 3Com without an RMA number will be returned to the sender unopened, at the sender's expense.

To obtain an RMA number, call or fax:

| Country                          | Telephone Number         | Fax Number         |
|----------------------------------|--------------------------|--------------------|
| U.S. and Canada                  | (800) 876 3266, option 2 | (408) 764 7120     |
| Europe                           | 31 30 60 29900, option 5 | (44) (1442) 275822 |
| Outside Europe, U.S., and Canada | (1) (408) 492 1790       | (1) (408) 764 7290 |

# INDEX

---

## Numerics

3Com Bulletin Board Service (3ComBBS) A-1  
3Com sales offices A-4  
3ComFacts A-3

---

## A

AARP 7-10  
AARP cache  
    administering 12-7  
    displaying 12-8  
    removing an entry from 12-9  
address  
    classes 4-3  
    IP to MAC, translating 9-13  
    MAC 3-3  
    network 3-3  
Address Resolution Protocol. *See* ARP  
Administration Console  
    menu descriptions 1-2  
    top-level menu 1-2  
ADSP 7-10  
advertisement address  
    adding 9-8  
    in a define interface command 9-5  
    removing 9-8  
AEP 7-8  
alarm thresholds  
    examples of 13-7  
    setting 13-7  
AppleTalk  
    address resolution protocol (AARP) 7-10  
    checksum 12-12  
    configuring forwarding 12-11  
    data stream protocol (ADSP) 7-10  
    echo protocol (AEP) 7-8  
    interface, displaying 12-3  
    main menu 1-6  
    name binding protocol (NBP) 7-9  
    network layer 7-6  
    nodes 7-2  
    physical layer 7-5  
    printer access protocol (PAP) 7-10

    protocols, about 7-1  
    protocols, and OSI levels 7-4  
    routing table maintenance protocol (RTMP) 7-6  
    routing tables 7-8  
    session layer protocol (ASP) 7-10  
    statistics, viewing 12-13  
    transaction protocol (ATP) 7-9  
    zone information protocol (ZIP) 7-9  
    zones 7-3  
AppleTalk networks 7-2  
    extended 7-2  
    nonextended 7-2  
AppleTalk node  
    pinging an 12-12  
AppleTalk routing 7-1  
ARP  
    defined 4-7, 9-13  
    location in OSI reference model 4-1  
    reply 4-8  
    request 4-8  
    *See also* ARP cache 9-13  
ARP cache 4-7, 9-13  
    flushing 9-15  
    removing an entry from 9-14, 9-19, 9-20  
ASP 7-10  
ATM ARP cache  
    displaying 9-17  
    flushing 9-18  
    removing an entry 9-17  
ATM ARP servers  
    about 4-10  
    defining 9-15  
    nodes that can function as an 4-10  
ATP 7-9

---

## B

BOOTP relay threshold 9-20  
bridge  
    menu 1-4  
bridge menu 1-3

bridging/routing  
     LANplex model 3-4  
     traditional model 3-4  
 bulletin board service A-1

---

## C

cache  
     displaying the IP multicast 10-9  
 checksum  
     configuring AppleTalk 12-12  
 chooser, Macintosh 7-2  
 CompuServe A-2  
 conventions  
     notice icons 2

---

## D

datagram delivery protocol 7-6  
 datagrams, statistics 9-23  
 data-link layer 4-1  
 DDP statistics 12-13  
 default route  
     displayed 9-11  
 default route, IP  
     defined 4-7, 9-10  
     removing 9-13  
     setting 9-13  
 defining 9-4  
 defining an ATM ARP 9-15  
     ATM ARP server 9-4  
 direct, route status 9-10  
 DVMRP  
     about 5-2  
     enabling 10-2  
     metric value 5-5, 10-3  
 dynamic routes 4-6, 6-14  
     *See also* RIP  
     *See also* SAP  
 dynamic routes, IPX 6-9

---

## E

extended network numbers 7-2  
 extended switching, overview 1-1

---

## F

fax service. *See* 3ComFacts  
 flushing  
     ARP cache 9-15

learned routes, AppleTalk 12-6  
 learned routes, IP 9-12  
 learned routes, IPX 11-7  
 for 11-9  
 forwarding  
     configuring AppleTalk 12-11

---

## G

gateway  
     routing table, and the 4-5  
     *See also* router

---

## H

hysteresis mechanism 13-8

---

## I

ICMP  
     defined 4-9  
     echo (request and reply) 9-22  
     echo Reply 4-9  
     echo request 4-9  
     ping and 9-22  
     Redirect 4-9  
     Time Exceeded 4-9  
 ICMP Router Discovery, enabling 9-21  
 IGMP  
     about 5-1  
     enabling 10-2  
 interface  
     administering an IP multicast 10-3  
     defining an IP 9-6  
     defining an IP multicast 10-2  
 interface, AppleTalk  
     defining an 12-3  
     displaying an 12-3  
     removing an 12-4  
 interface, IP  
     defining a LIS 9-4  
     defining a VLAN 9-6  
     displaying an 9-3  
     parts of 9-1  
     removing definition 9-7  
 interface, IP multicast  
     disabling 10-5  
     displaying 10-4  
     enabling 10-5  
     parts of 10-1  
 interface, IPX  
     defining an 11-3

- displaying an 11-3
- modifying an 11-4
- removing an 11-4
- Interior Gateway Protocols (IGP) 4-6, 6-9
- Internet address. *See* IP address
- Internet Control Message Protocol. *See* ICMP
- Internet Protocol. *See* references with IP address
- intranetwork routing
  - diagram 3-2
- IP
  - address translation 9-13
  - ARP cache 9-13
  - enabling routing 9-20
  - interface 9-1
  - main menu 1-4
  - pinging a station 9-22
  - RIP mode 9-21, 9-22
  - route
    - displaying table 9-11
  - routes 9-9
  - statistics, displaying 9-23
- IP address
  - classes of 4-3
  - defined 4-2
  - derived from 4-2
  - division of network and host 4-2
  - example 4-4
  - network layer and the 4-1
  - RIP, and 4-6
  - routing table, and the 4-5
  - subnet mask, and the 4-3
  - subnet part 4-3
- IP interface
  - defining 9-6
  - displaying 9-3
  - removing definition 9-7
- IP multicast routing
  - interface
    - disabling 10-5
- IP multicast
  - cache, displaying 10-9
  - routes, displaying 10-8
- IP Multicast menu 1-4
- IP multicast routing
  - about 5-1
  - algorithms 5-3
  - interfaces 10-1
    - administering 10-3
    - defining 10-2
    - displaying 10-4
    - enabling 10-5
- MBONE 5-2
- rate limit 5-6, 10-4

- TTL threshold 10-3
- tunnels 5-6, 10-6
- IP route
  - default 9-10, 9-13
  - defining static 9-11
  - removing from table 9-12
  - status 9-10
- IP router
  - transmission process 4-2
- IP routing
  - address classes 4-3
  - basic elements 4-2
  - enabling 9-20
  - ICMP 4-9
  - OSI reference model 4-1
  - references 4-11
  - router interface 4-4
  - routing table 4-5
  - transmission errors 4-9
- IP routing over ATM 4-10
- IPX
  - forwarding statistics, displaying 11-17
  - main menu 1-5
  - RIP statistics, displaying 11-15
  - route
    - defining a static 11-6
    - removing a 11-7
  - SAP statistics 11-16
  - static server 11-9
- IPX routing
  - and RIP 6-10
  - packet format 6-5
  - router interface 6-8
  - routing table 6-8
  - SAP, and 6-10
  - server table 6-13

---

## L

- LANplex
  - bridging/routing model 3-6
  - intranetwork router, as an 3-2
  - ports and IP interfaces 9-6
  - subnetting with 3-2
- learned routes
  - flushing AppleTalk 12-6
  - flushing IP 9-12
  - flushing IPX 11-7
- learned, IP route status 9-10
- LIS
  - definition of 4-10
  - forwarding to nodes within an 4-11

LIS interfaces  
     characteristics of 9-3  
     defining 9-4

## M

MAC (Media Access Control). *See* FDDI MAC  
 MAC address 3-3  
     ARP and 9-13  
     bridging in switching modules, and 3-6  
     compared to IP address 4-2  
     in ARP Request 4-8  
     located with ARP 4-7  
     use in IP routing 4-8  
 Macintosh, chooser 7-2  
 management  
     IP interface 9-1  
 management console  
     RMON 13-1  
 MBONE 5-2  
 menu  
     AppleTalk main 1-6  
     bridge 1-4  
     IP main 1-4  
     IPX main 1-5  
 metric  
     defined 4-5  
 metric value  
     DVMRP 5-5, 10-3  
 MIB  
     RMON 13-1, 13-2, 13-4  
 multicast routing, IP  
     about 5-1

## N

name binding protocol 7-9  
 named entities 7-2  
 NBP 7-9  
 NetWare  
     defined 6-1  
     OSI reference model, and the 6-2  
     protocols 6-1 to 6-3  
 network address 3-3  
 network layer, and IP address 4-1  
 network layer, AppleTalk 7-6  
 network numbers  
     extended 7-2  
     nonextended 7-2  
 network supplier support A-3  
 nodes, AppleTalk 7-2  
 nonextended network numbers 7-2

## O

on-line technical services A-1  
 OSI Reference Model  
     AppleTalk routing and 7-5  
     IP routing and 4-1  
     IPX routing and 6-2

## P

PAP 7-10  
 physical layer, AppleTalk 7-5  
 ping  
     AppleTalk node 12-12  
     IP station 9-22  
 port  
     *See also* FDDI port  
 printer access protocol 7-10  
 probe  
     RMON 13-1, 13-2  
 PVC  
     adding 9-9  
     removing 9-9

## R

rate limit  
     IP multicast 5-6, 10-4  
 references  
     Comer 4-11  
     Perlman 4-11  
     routing RFCs 4-11  
 returning products for repair A-4  
 RIP  
     active mode 9-21  
     broadcast address, and 9-2  
     default mode 9-22  
     defined 4-6, 6-10  
     off mode 9-21  
     passive mode 9-21  
     route configuration, and 4-6, 6-9  
     setting mode 9-21  
     using for dynamic routes 6-9  
 RIP statistics  
     IPX RIP 11-15  
 RMON  
     agents 13-3  
     alarms 13-6  
     benefits of 13-2  
     groups 13-3  
     hysteresis mechanism 13-8  
     LANplex implementation 13-2



- management console 13-1
- MIB 13-1, 13-2, 13-4
- probe 13-1, 13-2
- route, IP
  - default 9-10
  - defining static 9-11
  - removing default 9-13
  - removing from table 9-12
  - status 9-10
- route, IPX
  - removing a 11-7
- router interface, IP
  - described 4-4
  - diagram 4-5
  - routing table, and the 4-5
- router interface, IPX
  - described 6-8
- routers, seed 7-4
- routes, displaying IP multicast 10-8
- routing
  - and bridging in switching modules 3-4
  - and bridging, traditional model 3-4
  - implementation in LANplex 3-4
  - LANplex system, and the 3-1 to 3-7
  - See also* IP routing, IPX routing, and AppleTalk routing
- Routing Information Protocol. *See* RIP
- routing table
  - display routes 9-11
- routing table, AppleTalk 7-8
- routing table, IP
  - contents 4-5, 9-9
  - default route 4-7
  - default route, setting 9-13
  - described 4-5
  - dynamic routes 4-6
  - example 4-6
  - flushing learned routes 9-12
  - metric 4-5
  - removing default route 9-13
  - removing route 9-12
  - static routes 4-6
- routing table, IPX
  - contents 6-8
  - described 6-8
  - displaying 11-6
  - dynamic routes 6-9
  - example 6-9
  - flushing learned routes 11-7
  - removing a route 11-7
  - static routes 6-9
- RTMP
  - description of 7-6

---

## S

- SAP
  - aging mechanism 6-14
  - packet structure 6-11
  - request handling 6-15
  - using for dynamic routes 6-14
- SAP mode
  - setting 11-13
- SAP statistics, displaying 11-16
- seed routers 7-4
- segmentation, increasing 3-3
- server 9-4, 9-16
  - defining a static IPX 11-9
- server table
  - contents 6-13
  - described 6-13
  - displaying 11-9
- Service Advertisement Protocol. *See* SAP
- session layer protocols
  - AppleTalk 7-9
- software
  - installation 1-1
- static route, IP 4-6
  - status of 9-10
- static route, IPX 6-9
  - defining 11-6
- static server, IPX
  - defining a 11-9
- station
  - See also* FDDI station
- statistics
  - AppleTalk, viewing 12-13
  - IP 9-23
  - IPX forwarding 11-17
  - IPX SAP 11-16
  - ZIP, displaying 12-15
- subnet mask
  - defined 4-3
  - diagram 4-4
  - example 4-4
  - in routing table 4-5
- subnetting
  - defined 4-3
  - Ethernet switching and 3-2
  - subnet mask, and the 4-3
  - with the LANplex 3-2

---

## T

- technical support A-1
- ThreeComForum A-2

timing out, IP route status 9-10  
T-notify  
    configuring 8-4  
transmission errors  
    ICMP Redirect 4-9  
    reasons for 4-9  
TTL threshold 5-5  
    IP multicast 10-3  
tunnels  
    IP multicast 5-6, 10-6

---

## V

VLAN  
    information  
        defining 8-3  
        displaying 8-1  
        modifying 8-4  
        removing 8-5  
VLAN interfaces  
    about 9-1  
    characteristics of 9-2  
    defining 9-6  
VLANs  
    application oriented 2-2  
    MAC address group 2-2  
    overlapped IP 2-7  
    port group 2-1  
    protocol-sensitive 2-2  
    routing between 2-8

---

## Z

ZIP 7-9  
    statistics, displaying 12-15  
zone information protocol (ZIP) 7-9  
zone information table (ZIT) 7-9  
    displaying the 12-10  
zone, AppleTalk  
    default 12-3  
    example of 7-3  
    naming 12-3